FACULTY
OF MATHEMATICS
AND PHYSICS
Charles University

DOCTORAL THESIS

Vojtěch Kaluža

Metric and analytic methods

Department of Applied Mathematics

Supervisor of the doctoral thesis:	doc. RNDr. Martin Tancer, Ph.D.
Co-advisor of the doctoral thesis:	Univ.-Prof. Dr. Eva Kopecká
Study programme:	Computer Science
Study branch:	Discrete Models and Algorithms

Prague 2018

I declare that I carried out this doctoral thesis independently, and only with the cited sources, literature and other professional sources.

I understand that my work relates to the rights and obligations under the Act No. 121/2000 Sb., the Copyright Act, as amended, in particular the fact that the Charles University has the right to conclude a license agreement on the use of this work as a school work pursuant to Section 60 subsection 1 of the Copyright Act.

In .... date .... signature of the author

First of all, I would like to thank three people responsible for my scientific leadership. To Jiří Matoušek for introducing Feige's question to me and for sharing his intuition about the question with me. To Eva Kopecká for being a very generous host and for her insight which aimed us in the right direction when we were stuck at the beginning and which often helped to clarify our arguments. To Martin Tancer for many research opportunities he has mediated for me and for all the time he has spent thinking about open problems with me. I learnt a lot whenever I could see the way he approached an open question.

I would also like to thank all my collaborators, Michael Dymond, Eva Kopecká, Martin Tancer, Pavel Paták, Zuzana Patáková, Éric Colin de Verdière, Xavier Goaoc, Alfredo Hubard, Arnaud de Mesmay, Robert Šámal and Radek Hušek. Thank you for all the stimulating discussions and many wonderful ideas you provided.

My doctoral studies would not be possible without the financial support provided to me mostly by the Czech Science Foundation (GAČR), the Austrian Science Fund (FWF), the Charles University Grant Agency (GAUK) and also by several other occasional sources. I am also very grateful to all the people that made my connection to some of the funds possible, namely to Martin Tancer, Vít Jelínek, Eva Kopecká and to senior members of my department which have been trying to improve the financial support of graduate students as much as possible.

Last but not least, I express my deepest gratitude to my family, most importantly to my wife Adéla and to my parents. I am what I am thanks to you.

Title: Metric and analytic methods

Author: Vojtěch Kaluža

Department: Department of Applied Mathematics

Supervisor: doc. RNDr. Martin Tancer, Ph.D., Department of Applied Mathematics

Co-advisor: Univ.-Prof. Dr. Eva Kopecká, Institut für Mathematik, Universität Innsbruck

Abstract: The thesis deals with two separate problems. In the first part we show that the regular $n \times n$ grid of points in $\mathbb{Z}^2$ cannot be recovered from an arbitrary n^2-element subset of $\mathbb{Z}^2$ using only mappings with prescribed maximum stretch independent of n. This provides a negative answer to a question of Uriel Feige from 2002. The present approach builds on the work of Burago and Kleiner and McMullen from 1998 on bilipschitz non-realisable densities and bilipschitz non-equivalence of separated nets in the plane. We describe a procedure that takes a positive, measurable function and encodes it into a sequence of discrete sets. Then we show that applying this procedure to a typical positive, continuous function on the unit square yields a counter-example to Feige's question. Along the way we provide a new proof of a result on bilipschitz decomposition for Lipschitz regular mappings, which was originally proved by Bonk and Kleiner in 2002.

In the second part we provide a constructive proof for the strong Hanani–Tutte theorem on the projective plane. In contrast to the previous proof by Pelsmajer, Schaefer and Stasi from 2009, the presented approach does not rely on characterisation of embeddability into the projective plane via forbidden minors.

Keywords: non-realisable function, Lipschitz, bilipschitz, grid points, Feige, regular mapping, Hanani–Tutte, projective plane

Contents

List of symbols

$:=$	An equality used as a definition.	*page 16*

Symbols used in Chapter 1

$[n]$	The set $\{1, \dots, n\}$.	*page 16*
I	The unit interval $[0, 1]$.	*page 16*
$\overline{B}(A, r)$	The points at distance at most r from the set A.	*page 16*
$\overline{B}(x, r)$	The closed ball of radius r centred at x.	*page 16*
$B(A, r)$	The points at distance less than r from the set A.	*page 16*
$B(x, r)$	The open ball of radius r centred at x.	*page 16*
$\overline{A}$	The topological closure of a set A.	*page 16*
$\operatorname{int} A$	The topological interior of a set A.	*page 16*
∂A	The topological boundary of a set A.	*page 16*
$\operatorname{diam}(A)$	The Euclidean diameter of a set A.	*page 16*
$\operatorname{dist}(x, A)$	The Euclidean distance of a point x to a set A.	*page 16*
$\\|u\\|_2$	The Euclidean norm of a vector u.	*page 16*
$\\|u\\|_\infty$	The supremum norm of a vector u.	*page 16*
$\\|L\\|_{\mathrm{op}}$	The operator norm of a linear operator L.	*page 16*
$\nu_i \rightharpoonup \nu$	The weak convergence of measures ν_i to measure ν.	*page 17*
$\rho\mathcal{L}$	The measure with density ρ with respect to the Lebesgue measure.	*page 17*
$f_\sharp\nu$	The pushforward measure of a measure ν induced by a mapping f.	*page 17*
$\fint_A \rho$	The average value of a function ρ over a set A, i.e., $\dfrac{1}{\mathcal{L}(A)} \displaystyle\int_A \rho \, \mathrm{d}\mathcal{L}$.	*page 17*
$\rho\|_C$	The restriction of a mapping ρ or a measure ρ to a set C.	*page 17*
f^{-1}	The inverse mapping to f.	*page 17*

$f^{-1}(B)$	The preimage of a set B under a mapping f.	*page 17*
$\mathrm{Jac}(f)(x)$	The Jacobian of a mapping f at a point x, i.e., $\det(Df(x))$.	*page 18*
$Df(x)$	The (Fréchet) derivative of a mapping f at a point x.	*page 18*
$C^1(A;\mathbb{R}^n)$	The space of continuously differentiable mappings $A \to \mathbb{R}^n$ with the supremum norm.	*page 18*
$C(A;\mathbb{R}^n)$	The space of continuous mappings $A \to \mathbb{R}^n$ with the supremum norm.	*page 18*
$C(A)$	The space of continuous functions $A \to \mathbb{R}$ with the supremum norm.	*page 18*
$\lVert f\rVert_{L^\infty}$	The essential supremum of a function $\lvert f\rvert$.	*page 18*
$L^\infty(A)$	The space of essentially bounded measurable functions on a set A.	*page 18*
(C,L)-regular	A Lipschitz regular mapping with Lipschitz constant at most L and $\mathrm{Reg}(f) \le C$.	*page 19*
$\mathrm{Reg}(f)$	The regularity constant of a Lipschitz regular mapping f; see Definition 1.12.	*page 19*
$N(f)$	The set of 'non-critical points' of a mapping f; see Definition 1.24.	*page 33*
$\mathbf{e}_1$-adjacent	Two cubes S, S' of side λ are $\mathbf{e}_1$-adjacent if $S' = S + \lambda\mathbf{e}_1$.	*page 46*
$\mathbf{e}_i$	The i-th canonical vector in $\mathbb{R}^d$, i.e., the unit vector of which all but i-th co-ordinates are 0.	*page 46*

Symbols used in Chapter 2

$G - e$	The graph obtained from G by removing the edge e.	*page 80*
$G - F$	The graph obtained from G by removing the edges in F, where $F \subseteq E(G)$.	*page 80*
$G - v$	The graph obtained from G by removing the vertex v and all its adjacent edges.	*page 80*
$G - W$	The graph obtained from G by removing the vertices in W, where $W \subseteq V(G)$.	*page 80*
$\otimes$	The crosscap.	*page 81*

$\mathbb{R}P^2$ — The projective plane. We realise it as the sphere with a crosscap. *page 81*

S^2 — The 2-dimensional sphere. *page 81*

$\mathrm{cr}_D(e, f)$ — The number of crossings between the edges e and f in the drawing D counted modulo 2. *page 83*

$D(\omega)$ — Let D be a drawing of a graph G and ω be a walk in G. Then $D(\omega)$ stands for the curve defined as the concatenation of the drawings of the edges of ω in the order and the directions determined by ω. *page 83*

$D(H)$ — Let D be a drawing of a graph G and H be a subgraph of G. Then $D(H)$ stands for the part of the drawing D corresponding to H. *page 83*

HT-drawing — A Hanani–Tutte drawing of a graph. *page 82*

(D, λ) — A projective HT-drawing D of a graph on S^2. *page 85*

$\lambda(\omega)$ — The type of a walk ω in a projective HT-drawing of a graph on S^2. *page 85*

S^+ — The component of S^2 corresponding to the inside of a separating cycle. *page 89*

S^- — The component of S^2 corresponding to the outside of a separating cycle. *page 89*

V^+, E^+ — The inside vertices and the inside edges, respectively. *page 89*

V^-, E^- — The outside vertices and the outside edges, respectively. *page 89*

G^{+0} — The part of the graph G that consists of the inside vertices and edges together with the vertices and the edges of a separating cycle. *page 89*

G^{-0} — The part of the graph G that consists of the outside vertices and edges together with the vertices and the edges of a separating cycle. *page 89*

A^+ — The inside arrow (multi-)graph. *page 90*

A^- — The outside arrow (multi-)graph. *page 90*

$\overline{uv}$ — The arrow between the vertices u and v. *page 91*

W^+_{uv} — The set of all proper, non-trivial walks in G^{+0} with endpoints u and v. *page 91*

W_{uv}^{-}	The set of all proper, non-trivial walks in G^{-0} with endpoints u and v. *page 91*
$(\mathbb{R}P^2)^+$	The part of the projective plane corresponding to the inside of a separating cycle. *page 107*
$(\mathbb{R}P^2)^-$	The part of the projective plane corresponding to the outside of a separating cycle. *page 107*
L_B	A valid labelling for the bridge B. *page 103*

Preface

In autumn 2011 I was a third-year undergraduate student of computer science looking for an advisor and a topic for bachelor thesis. I had no idea what I would like to work on. I just knew that during my studies I generally enjoyed mathematical or theoretical subjects more than those concerned with programming or various kinds of software and hardware, my favourite being Mathematical analysis. With these feelings, I addressed Jiří Matoušek, who then suggested, as a possible topic for my thesis, a question asked by Uriel Feige. The question concerns finite sets of points in the plane with integer co-ordinates which one is then asked to map bijectively onto a regular square grid of points in such a way that the pairwise distances between the points are stretched as little as possible. Seen more abstractly, the question is about embedding finite metric spaces coming from discrete sets in the plane into discrete metric spaces of a very special form with as low distortion as possible. Feige asked the question in 2002 during a workshop held in Haifa that was aimed at this type of problems, which was also attended by Matoušek.

When presenting the problem to me, Matoušek said that the problem was hard in his opinion. So my initial goal was not to solve the problem, but to get acquainted with it, to study related literature, and possibly, to try to solve some of its special cases. Matoušek pointed me to work on equivalence of separated nets and prescribed Jacobian equation by Dmitri Burago, Bruce Kleiner and Curtis McMullen and said that maybe similar methods could be employed to attack Feige's question.

Since then, the question has become my faithful companion backing me also through my master thesis to the beginning of my doctoral studies, all done under the supervision of Matoušek.

However, around the autumn 2014, when I became a first-year graduate student, Matoušek's health started to deteriorate. Nevertheless, he contacted Eva Kopecká from Universität Innsbruck, who works mainly in Lipschitz analysis and who then invited me to come to Innsbruck. Later, she hosted me in Innsbruck for a number of longer and shorter stays, and also, became my co-advisor.

This way I teamed up with Eva Kopecká and Michael Dymond, a post-doc in Innsbruck. We worked together on Feige's question for more than a year until we finally managed to resolve it. Our findings constitute the core of the presented thesis.

Sadly, Matoušek passed away at the beginning of March 2015. I grouped with Kopecká and Dymond only several months after that. To answer Feige's question, we really used tools heavily inspired by the work of Burago, Kleiner and McMullen, as was suggested to me by Matoušek at the beginning of the story. Thus, at the highest level, this idea should be attributed to him. However, I would like to note that in my work towards answering Feige's question under Matoušek's supervision I did not really move beyond understanding the related work of Burago, Kleiner and McMullen. All the new results that led to the resolution of Feige's question arose in collaboration with Dymond and Kopecká only later.

After Matoušek's passing, I had to find a new supervisor. Martin Tancer was willing to take the role. His main interests lie in topological and geometric combinatorics and computational topology, which naturally led to extension of my own interest towards these areas.

Martin Tancer has allowed me to connect with several of his foreign collaborators. He regularly invited me to join when he was hosting visitors for short research stays and also often enabled me to accompany him when he was visiting one or more of his colleagues abroad. One of the problems that I started to work on during the visits, together with Éric Colin de Verdière, Pavel Paták, Zuzana Patáková, Martin Tancer and, at the beginning, with Alfredo Hubard, was a possible extension of the strong Hanani–Tutte theorem to surfaces other than the plane. We have not quite succeeded in that, however, we were at least able to provide a constructive proof of the strong Hanani–Tutte theorem for the projective plane. This proof forms the second part of the present thesis.

The content of Chapter 1, which discusses Feige's question, is based on the article [Paper I] written in collaboration with Dymond and Kopecká, which is accepted for publication in *Geometric and Functional Analysis*.

The content of Chapter 2, which is about the strong Hanani–Tutte theorem for the projective plane, is based on the article [Paper II] written in collaboration with Colin de Verdière, Paták, Patáková and Tancer, which is published in Journal of Graph Algorithms and Applications.

At the end, I would like to explain the way I use pronouns in the thesis. I follow widespread habit of using 'we' in the sense 'the author' together with 'the kind reader'. However, since the work presented in the thesis was done in collaboration, I wanted to make clear to the reader that I am not the only author of the results. Therefore, at several places I refer to myself and my co-authors using 'they'. This is not to dispose myself of responsibility, it is solely intended to give a proper credit to all people involved. Whenever I express my personal views, I refer to myself as 'the author'.

List of author's publications used in the thesis

[Paper I] Michael Dymond, Vojtěch Kaluža and Eva Kopecká
2018. Mapping n grid points onto a square forces an arbitrarily large Lipschitz constant. *Geometric and Functional Analysis*.
`http://dx.doi.org/10.1007/s00039-018-0445-z`

[Paper II] Éric Colin de Verdière, Vojtěch Kaluža, Pavel Paták, Zuzana Patáková and Martin Tancer
2017. A Direct Proof of the Strong Hanani–Tutte Theorem on the Projective Plane. *Journal of Graph Algorithms and Applications*, 21(5):939–981.
`http://dx.doi.org/10.7155/jgaa.00445`

1. Feige's question

Introduction

History and motivation

The motivation for the work presented in this chapter stems out of the so-called **Graph Bandwidth problem**, which is defined in the following way.

Consider a graph $G = (V, E)$ with n vertices. Assume we are given a bijection $f\colon V \to \{1, \dots, n\}$. Let us call the quantity $|f(u) - f(v)|$ a *stretch* of an edge (u, v) under f. The minimum taken over all bijections f as above of maximal stretch of an edge under f is called the *bandwidth* of a graph. The goal in the Graph Bandwidth problem is to find an f realising the bandwidth of G.

The Graph Bandwidth problem can be stated using matrices: we say that an $n \times n$ matrix A has a *bandwidth* k if all its non-zero entries are concentrated in a band of width k along the main diagonal. Then the **Matrix Bandwidth problem** is a question of finding an $n \times n$ permutation matrix Π such that $\Pi A \Pi^T$ has the smallest possible bandwidth.

It is easy to see that the Graph bandwidth problem and the Matrix bandwidth problem for *symmetric* matrices are equivalent—the correspondence being given by the edges (i, j) and the non-zero entries $A_{i,j}$. Matrices of low bandwidth are easier to store and to manipulate with. Therefore, both bandwidth problems have been studied since the development of computers: according to Chinn, Chvátalová, Dewdney, and Gibbs [9], the Matrix Bandwidth problem has been considered at least since 1950s. The Graph Bandwidth problem originated, probably independently, in 1962 in jet Propulsion Laboratory in Pasadena [9]. Since then there has been an instantly growing literature on the two problems. As of March 2018, Google Scholar lists more than 400 articles containing the phrase 'Graph Bandwidth' and more than 2100 articles containing the phrase 'Matrix Bandwidth'. More recent surveys than [9] were written by Díaz, Petit, and Serna [17] and Mafteiu-Scai [36]. The latter focuses only on the practical point of view.

In 1976 Papadimitriou [44] showed that the decision version of the Graph Bandwidth problem is **NP**-hard. Much later, in 1998 Blache, Karpinski, and Wirtgen [4] showed that the Graph Bandwidth is **NP**-hard to approximate with a factor better than $\frac{3}{2}$ on general graphs and $\frac{4}{3}$ on trees. Unger [46] announced that it is **NP**-hard to approximate the Graph Bandwidth by any constant factor, even for caterpillar trees. However, the full proofs appeared only much later in the work of Dubey, Feige, and Unger [18].

On the positive side, Feige [23] designed a randomised algorithm that produces a linear layout of an n-vertex graph, i.e., a bijection between the set of its vertices and the set $\{1, \dots, n\}$, with a stretch of an edge at most $O(\log^{3.5} n\sqrt{\log\log n})$-times the optimal bandwidth. Feige's algorithm falls into a class of *approximate metric embeddings algorithms*; nowadays a rich area of research initiated arguably by a work of Bourgain [6]. Feige's algorithm was inspired by a famous work of Linial, London, and Rabinovich [35], who, in turn, built upon the previous work of Bourgain [6] and presented an algorithm constructing a *low-distortion embedding* of a weighted graph into a Euclidean space (together with numerous applications).

Feige's algorithm can be described, in very rough terms, in two steps. First, construct a random injection of the set of vertices of an n-vertex graph into a Euclidean space of dimension $O\left(\mathrm{polylog}\left(n\right)\right)$ that preserves the shortest-path metric[1] on the graph up to a factor of $O(\text{polylog}(n))$. Second, take a random line in the space and project the points representing the vertices of the graph orthogonally onto that line. The linear arrangement of vertices along the line then provides, with high probability, the desired approximate solution to the Graph Bandwidth problem.

Besides looking for a linear layout of the vertices of a graph, we can ask what is the best arrangement of the vertices of an n^2-vertex graph onto a grid $\{1,\dots,n\} \times \{1,\dots,n\}$ (or, analogously, for higher-dimensional grids). This is a natural generalisation of the Graph Bandwidth problem. Feige [private communication] realised that his algorithm would work just fine even for the generalised problem, except possibly the last step: instead of taking a linear arrangement along a line, one is dealing with an embedding into a plane (or a higher-dimensional subspace) that is then rounded to the grid points. However, it is not clear at all how to map the resulting set of points injectively onto $\{1,\dots,n\} \times \{1,\dots,n\}$ without increasing the stretch of the edges too much, because one ends up with a placement of the vertices in a larger grid that contains 'holes'. Motivated by this, Feige asked the following question:

Question 1.1 (Feige's question)**.** *Is there a universal constant $L > 1$ such that for every $n \in \mathbb{N}$ and every set $S \subset \mathbb{Z}^2$ of n^2 points there is an L-Lipschitz bijection between S and the square grid $\{1,\dots,n\} \times \{1,\dots,n\}$?*

It seems that the question has first received wider attention at the workshop 'Discrete Metric Spaces and their Algorithmic Applications' held in Haifa in 2002; see [40]. It also appeared in a list of open problems maintained by Matoušek and Naor [39] as Question 2.12.

We resolve Question 1.1 to the negative. The results presented in this chapter are contained in the work by Dymond, the author, and Kopecká [21].

To the best of author's knowledge, there has been almost no progress towards answering Question 1.1 before the presented work; the only prior published works considering the question being author's Master's [33] and Bachelor's theses [32], which, however, contain only a few minor results for very special cases.

At the highest level, the presented solution works as follows. On the one hand, we translate Question 1.1, which is a discrete problem, into a question in a continuous world, following the lines of a similar reduction in [8]. More specifically, we show that if the answer to Question 1.1 is positive, then every density of measure can be *realised* (a notion to be defined precisely later) as a volume form transformation induced by a special Lipschitz mapping (called Lipschitz regular). On the other hand, we show that not all densities can be realised in such a way. Two key steps in this part are an inspection and a strengthening of the construction of bilipschitz non-realisable functions by Burago and Kleiner [8] and a bilipschitz decomposition result for Lipschitz regular mappings following from a result by Bonk and Kleiner [5].

[1] In fact, it is not sufficient to take care of distances of pairs of vertices, i.e., to have a low-distortion embedding, but to control the behaviour of larger subsets of vertices. For this, Feige devised a so-called *volume respecting embedding*.

Interplay of discrete and continuous worlds

Burago and Kleiner [8] and, independently, McMullen [42] considered the following question, at that time open, that was first mentioned by Furstenberg in the 1960s and appears in the book [27] by Gromov.

Question 1.2. *Is there for every separated net*[2] *S in $\mathbb{R}^2$ a bilipschitz bijection $f\colon S \to \mathbb{Z}^2$. In other words, are every two separated nets in $\mathbb{R}^2$ bilipschitz equivalent?*

The approach of Burago and Kleiner [8] and McMullen [42] to answer Question 1.2 is based on its relation to another question, at that time open as well, which is suited in the continuous world:

Question 1.3. *Is every measurable function $\rho\colon [0,1]^2 \to (0,\infty)$ such that $0 < \inf \rho \le \sup \rho < \infty$* realisable *as a Jacobian of a bilipschitz mapping? That is, is there always a bilipschitz mapping $f\colon [0,1]^2 \to \mathbb{R}^2$ such that the equation*

$$\mathrm{Jac}(f) = \rho \tag{1.1}$$

holds almost everywhere?

Question 1.3 falls into the area of prescribed Jacobian of a homeomorphisms; for related work see Dacorogna and Moser [11] or Rivière and Ye [45], for example.

Burago and Kleiner [8] proved that a positive answer to Question 1.2 yields a positive answer to Question 1.3 (McMullen [42] established the opposite implication as well; consequently, the problems in Questions 1.2 and 1.3 are equivalent). Their proof presents a procedure that takes any function ρ as in the statement of Question 1.3 and encodes it into a separated net S in $\mathbb{R}^2$. Then Burago and Kleiner show that a bilipschitz bijection $S \to \mathbb{Z}^2$ can be used to obtain a bilipschitz homeomorphism $f\colon [0,1]^2 \to \mathbb{R}^2$ such that $\mathrm{Jac}(f) = \rho$ almost everywhere. As a next step, they construct a function ρ that cannot be realised as a Jacobian of a bilipschitz homeomorphism in the sense of equation (1.1), and thus, provide negative answers to both Questions 1.2 and 1.3. The work in [8] is stated in dimension $d = 2$ just for simplicity; it works without any significant change in any dimension $d \ge 2$.

The approach of Dymond, the author, and Kopecká [21] to Question 1.1 presented in this thesis follows the same pattern and works in any dimension $d \ge 2$ as well[3]. Consequently, a negative answer to the continuous analogue of Feige's question is provided in any dimension $d \ge 2$.

Question 1.4 (Feige's question for general dimension)**.** *Let $d \in \mathbb{N}$ be at least 2. Is there a universal constant $L := L(d)$ such that for every $n \in \mathbb{N}$ and every set $S \subset \mathbb{Z}^d$ of n^d points there is an L-Lipschitz bijection from S to the cubical grid $\{1, \dots, n\}^d$?*

We present a continuous question that is related to Feige's question. Before we state it, we note that in many of the statements below we use a few notions

[2]The set S is separated if no two distinct points of S are at distance less than some $R > 0$. It is a net of the ambient space if every point in the space is at distance at most $r > 0$ from S.

[3]In dimension $d = 1$ the answer to Feige's question is trivially positive.

and symbols that the reader may not be familiar with. All of them will be defined precisely in 'Background and notation'. In this section we provide only rough explanation of their intended meaning. To begin with, we write $\mathcal{L}$ for the Lebesgue measure.

Question 1.5 (A continuous analogue of Feige's question)**.** *Is there for every measurable function $\rho\colon [0,1]^d \to (0,\infty)$ such that $0 < \inf \rho \leq \sup \rho < \infty$ a Lipschitz regular[4] mapping $f\colon [0,1]^d \to \mathbb{R}^d$ verifying the equation*

$$\int_{f^{-1}(A)} \rho \, \mathrm{d}\mathcal{L} = \mathcal{L}(A) \qquad \textit{for every } A \subseteq f\left([0,1]^d\right)? \tag{1.2}$$

We prove that a positive answer to Question 1.4 provides a positive answer to the continuous question. It should be emphasized that we provide only one-sided relation between the two questions stated above, in contrast to the relation between Questions 1.2 and 1.3 established by Burago and Kleiner [8] and McMullen [42].

Theorem 1.6. *A positive answer to Question 1.4 yields a positive answer to Question 1.5.*

The proof of the theorem above is a modification of the discretisation argument of Burago and Kleiner [8] and is contained in Section 1.1.

As the second step towards answering Question 1.4, we show that there are many functions for which the answer to the continuous analogue of Feige's question is false. In fact, it is sufficient to look for such examples inside the space of continuous functions. Namely, we prove:

Theorem 1.46. *The set*

$$\mathcal{E} := \left\{ \rho \in C([0,1]^d) : \rho \textit{ admits a Lipschitz regular solution to equation } (1.2) \right\}$$

forms a σ-porous[5] subset of the space of continuous functions $C([0,1]^d)$ with the supremum norm.

We note that equation (1.2) is a generalisation of equation (1.1) from Question 1.3, and thus, Theorem 1.46 provides a twofold generalisation of the results of Burago and Kleiner [8] and McMullen [42]; first, a stronger notion of realisability taking into account Lipschitz regular mappings instead of bilipschitz is considered, and second, it is proven that there are many such non-realisable continuous functions instead of constructing just one example.

The proof of Theorem 1.46 and of necessary preliminary lemmas spans most of Sections 1.2, 1.3 and 1.4.

In Section 1.2 we will prove the following decomposition theorem, which says that any Lipschitz regular mapping decomposes into a finite number of bilipschitz mappings on some open subset of its image:

[4] Lipschitz regular mapping is a special kind of Lipschitz mapping that is in a sense non-degenerate: it cannot map a set of positive measure to a set of measure zero.

[5] It means that a typical continuous function is not in $\mathcal{E}$.

Theorem 1.35. *Let $U \subseteq \mathbb{R}^d$ be non-empty and open and $f\colon \overline{U} \to \mathbb{R}^d$ be a Lipschitz regular mapping. Then there exist a non-empty open set $T \subseteq f(\overline{U})$, $N \in \{1, \dots, \mathrm{Reg}(f)\}$ and pairwise disjoint open sets $W_1, \dots, W_N \subseteq \overline{U}$ such that $f^{-1}(T) = \bigcup_{i=1}^{N} W_i$ and $f|_{W_i}\colon W_i \to T$ is a bilipschitz homeomorphism for each i with the lower bilipschitz constant $b = b(\mathrm{Reg}(f))$.*

We note that, in the statement above, $\mathrm{Reg}(f)$ is a natural number that quantifies how 'regular' the Lipschitz regular mapping f is.

Theorem 1.35 is really crucial to overcome the gap between Lipschitz regular and bilipschitz mappings. It can be deduced quite easily from the following result of Bonk and Kleiner [5]. We present only a special case that we shall need:

Theorem 1.30 (Bonk and Kleiner [5])**.** *Let $U \subseteq \mathbb{R}^d$ be open and $f\colon \overline{U} \to \mathbb{R}^d$ be Lipschitz regular. Then there are disjoint open sets $(A_n)_{n\in\mathbb{N}}$ such that $\bigcup_{n\in\mathbb{N}} A_n$ is dense in $\overline{U}$ and such that $f|_{A_n}$ is bilipschitz with the lower bilipschitz constant $b = b(\mathrm{Reg}(f))$.*

When working on the first version of the presented results, Dymond, the author, and Kopecká [20] were not aware of the work by Bonk and Kleiner [5], and thus, devised and independent proof of Theorem 1.30. They were informed about [5] only later by Guy C. David. We will present a proof of Theorem 1.30 that is easier and shorter, though less general than that of Bonk and Kleiner in Section 1.2.

Then, in Section 1.3 we extract from the construction of Burago and Kleiner [8] a certain geometric property of Jacobians of bilipschitz mappings that is a key part of the construction of bilipschitz non-realisable functions. Adding several strengthenings not present in [8], we deduce the following statement:

Lemma 1.40. *Let $d, k \in \mathbb{N}$ with $d \geq 2$, $L \geq 1$ and $\eta, \zeta \in (0,1)$. Then there exists $r = r(d,k,L,\eta,\zeta) \in \mathbb{N}$ such that for every non-empty open set $U \subseteq \mathbb{R}^d$ there exist finite tiled families $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_r$ of cubes contained in U with the following properties:*

1. *For each $1 \leq i < r$ and each cube $S \in \mathcal{S}_i$*

$$\mathcal{L}\left(S \cap \bigcup_{j=i+1}^{r} \bigcup \mathcal{S}_j\right) \leq \eta \mathcal{L}(S).$$

2. *For any k-tuple $(h_1, \dots, h_k)$ of L-bilipschitz mappings $h_j\colon U \to \mathbb{R}^d$ there exist $i \in \{1, \dots, r\}$ and $\mathbf{e}_1$-adjacent cubes $S, S' \in \mathcal{S}_i$ such that*

$$\left| \frac{1}{\mathcal{L}(S)} \int_S |\mathrm{Jac}(h_j)| \, \mathrm{d}\mathcal{L} - \frac{1}{\mathcal{L}(S')} \int_{S'} |\mathrm{Jac}(h_j)| \, \mathrm{d}\mathcal{L} \right| \leq \zeta$$

   for all $j \in \{1, \dots, k\}$.

For a precise definition of the terms 'tiled' family of cubes and '$\mathbf{e}_1$-adjacent' cubes we refer the reader to Section 1.3. However, the rough meaning is that a 'tiled' family of cubes is formed by congruent cubes arranged into a tiling and two such cubes are '$\mathbf{e}_1$-adjacent' if they share the facet defined by fixing the first co-ordinate.

Finally, in Section 1.4 we show how to prove Theorem 1.46 using Lemma 1.40 and Theorem 1.35. We also show that a version of Theorem 1.46 for the space $L^\infty([0,1]^d)$ instead of $C([0,1]^d)$ holds as well, though providing only a weaker conclusion speaking about bilipschitz realisability only instead of the more general Lipschitz regular realisability.

Background and notation

In this Section we fix the notation used throughout the chapter and present the basic results that we shall use. Although many concepts that we use make sense in the setting of metric spaces, we stick here to the Euclidean spaces wherever possible since we will not work in any more general space.

We often write := to emphasize that the equality in question is to be interpreted as a definition. We write $[n]$ for the set $\{1,\dots,n\}$, where $n \in \mathbb{N}$. The unit interval $[0,1]$ is denoted by I.

Balls, norms, sets and nets

We denote an open *Euclidean ball* centred at $x \in \mathbb{R}^d$ with radius $r > 0$ by $B(x,r)$. In case the ambient space is not the Euclidean space, but a metric space (which will be the case only for spaces of functions), we mean by $B(x,r)$ an open metric ball of radius r centred at x. For a *closed ball* we write $\overline{B}(x,r)$. We also extend the notation to denote neighbourhoods of sets. Namely, $B(A,r)$ denotes the set of points at distance less than r to the set A (including the set A), and similarly, we define $\overline{B}(A,r)$.

For a general set $A \subseteq \mathbb{R}^n$ we denote by $\overline{A}, \operatorname{int} A, \partial A$ its topological *closure*, *interior* and *boundary*, respectively. For an introduction to topology, we refer the reader to classical monographs by Munkres [43] and Hatcher [28].

We write $\|\cdot\|_2$ for the *Euclidean norm.* For a set A and a point x we denote by $\operatorname{dist}(x,A)$ the *Euclidean distance* of x to A. Thus, $\operatorname{dist}(x,A) := \inf_{y\in A} \|x-y\|_2$. The *diameter* of a set A is defined as $\operatorname{diam}(A) := \sup_{x,y\in A} \|x-y\|_2$.

Sometimes we will also use the *supremum norm*, denoted by $\|\cdot\|_\infty$. The *operator norm* is denoted by $\|\cdot\|_{\mathrm{op}}$.

A set $S \subseteq \mathbb{R}^d$ is said to be *s-separated* if any two of its distinct points $x, y \in S$ are at distance $\|x-y\|_2 \geq s$. For $S \subseteq A$, we call S an *r-net* of A if any point of A is at distance at most r from S. We say that S is an *(s,r)-separated* net of A if it is an r-net of A which is s-separated. If we say that S is just a separated net, we mean that there exist $s, r > 0$ such that S is an (s,r)-net.

Measures

As a general reference for measure theory and the theory of the Lebesgue integral we suggest the books by Mattila [41] and Federer [22].

We adopt the convention of Mattila [41] and do not distinguish between *measures* and *outer measures.* We will encounter only *Borel measures*—measures for which every Borel set is measurable. Recall that a set is *Borel* if it can be generated from open sets using any combination of countable unions and countable intersections. For a Borel measure ν on $\mathbb{R}^d$ there is the unique smallest closed set F such that $\nu(\mathbb{R}^d \setminus F) = 0$. The set F is called the *support* of ν.

We denote the d-dimensional *Lebesgue measure* by $\mathcal{L}^d$. If the dimension of the ambient space is understood, we usually write just $\mathcal{L}$ instead of $\mathcal{L}^d$.

Let ν be a measure on a set A. For $B \subseteq A$ we denote by $\nu|_B$ the restriction of the measure ν to a set B, i.e., the measure defined as $\nu|_B(C) := \nu(B \cap C)$. To a ν-measurable mapping $f\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^n$ we associate a measure $f_\sharp\nu$ on

$f(A)$ via the formula $f_\sharp\nu(D) := \nu(f^{-1}(D))$ for every $D \subseteq f(A)$. We call $f_\sharp\nu$ the *f-pushforward of ν*, or just the *pushforward* of ν.

We say that a function $\rho\colon A \to [0,\infty)$ is a *density* of measure ν if it is a density of ν with respect to the Lebesgue measure, that is,

$$\int_B \rho \,\mathrm{d}\mathcal{L} = \nu(B)$$

for every $B \subseteq A$. For the measure with density ρ we write $\rho\mathcal{L}$.

Given $A \subseteq \mathbb{R}^d$ such that $0 < \mathcal{L}(A) < \infty$, by the *average value* of a function $\rho\colon A \to \mathbb{R}$ we mean the quantity

$$⨍_A \rho := \frac{1}{\mathcal{L}(A)} \int_A \rho \,\mathrm{d}\mathcal{L}.$$

We write that a property or a formula $P(x)$ holds *ν-a.e.* (meaning almost everywhere with respect to a measure ν) in A if the set $\{x \in A\colon \neg P(x)\}$ has ν measure zero. Whenever we do not specify the measure ν and use just 'a.e.', we mean $\mathcal{L}$-a.e.

By the *weak convergence* of finite Borel measures ν_i to a finite Borel measure ν, denoted by $\nu_i \rightharpoonup \nu$, we mean the convergence of $\int \varphi \,\mathrm{d}\nu_i$ to $\int \varphi \,\mathrm{d}\nu$ for every continuous function φ with compact support. Every sequence of measures has at most one weak limit.

Mappings, functions and derivatives

Let $f\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^n$ be an arbitrary mapping. By f^{-1} we denote its *inverse* whenever it exists. For a set $B \subseteq \mathbb{R}^n$, even in the case that f^{-1} does not exist, we write $f^{-1}(B)$ for the *preimage* of B under f, that is, for the set $\{x \in A\colon f(x) \in B\}$. For a set C we denote by $f|_C$ the *restriction* of a mapping f to a set C.

A mapping f is *Borel* if $f^{-1}(A)$ is Borel set for every open set A in the image of f. In particular, every continuous mapping is Borel. Moreover, every Borel mapping is Lebesgue measurable. We recall the general form of change of variables for Borel mappings, which can be found in Mattila [41, Thm. 1.19], for example:

Theorem 1.7 (Change of variables)**.** *Let $f\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^n$ and $g\colon f(A) \to [0,\infty)$ be Borel mappings. If ν is a Borel measure on A, then*

$$\int_A g \circ f \,\mathrm{d}\nu = \int_{f(A)} g \,\mathrm{d}f_\sharp\nu.$$

We cannot introduce here all concepts and theorems from topology that we are going to use. At the moment, we state only one topological tool, which we are going to use very often—the so-called 'Invariance of Domain' was proven by Brouwer [7] and says the following:

Theorem 1.8 (Invariance of Domain [7])**.** *Let $U \subseteq \mathbb{R}^d$ be open and $f\colon U \to \mathbb{R}^d$ be continuous and injective. Then $f(U)$ is open and f is a homeomorphism onto its image.*

When the dimension of the domain and the range of f agree, we write $Df(x)$ for the *Fréchet derivative* of f at x and $\mathrm{Jac}(f)(x) := \det(Df(x))$, called the *Jacobian* of f at x, for its determinant. Geometrically, $|\mathrm{Jac}(f)(x)|$ describes the factor by which the volume around x is locally changed by f and the sign of $\mathrm{Jac}(f)$ determines the change of orientation at x.

For a set $A \subseteq \mathbb{R}^d$ we write $C(A)$ for the *space of continuous functions* $A \to \mathbb{R}$ endowed with the supremum norm. Briefly, we will work with spaces of continuous mappings $A \to \mathbb{R}^n$ with the supremum norm, which we will denote by $C(A;\mathbb{R}^n)$, and the space of *continuously differentiable* mappings $A \to \mathbb{R}^n$ denoted by $C^1(A;\mathbb{R}^n)$.

At one place in Section 1.4, we will also use the space $L^\infty(A)$ of all *essentially bounded measurable* functions with respect to the Lebesgue measure endowed with the L^∞-norm $\|\cdot\|_{L^\infty}$, where

$$\|\rho\|_\infty := \inf\{C > 0\colon\ |\rho(x)| \leq C \text{ for a.e. } x \in A\}.$$

Lipschitz and bilipschitz mappings

Let $L > 0$. We say that $f\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^n$ is *L-Lipschitz* if

$$\|f(x) - f(y)\|_2 \leq L\,\|x - y\|_2$$

for every $x, y \in A$. Let $0 < b \leq L$. We say that f is (b, L)-*bilipschitz* if

$$b\,\|x - y\|_2 \leq \|f(x) - f(y)\|_2 \leq L\,\|x - y\|_2$$

for $x, y \in A$. Often we say that f is L-bilipschitz, which means that it is $\left(\frac{1}{L}, L\right)$-bilipschitz. In other words, f is L-bilipschitz if both f and f^{-1} are L-Lipschitz. We say that f is Lipschitz or bilipschitz, if there is $L > 0$ for which it is L-Lipschitz or L-bilipschitz, respectively. The smallest such L is called the *Lipschitz* or *bilipschitz constant* of f.

An immediate consequence of the definition of the Lebesgue measure is that an L-Lipschitz mapping $f\colon \mathbb{R}^d \to \mathbb{R}^d$ cannot increase the Lebesgue measure[6] of a set by a factor larger than L^d. In other words, for any $A \subseteq \mathbb{R}^d$ we have that

$$\mathcal{L}(f(A)) \leq L^d \mathcal{L}(A). \tag{1.3}$$

We state here two well-known and classical results on Lipschitz mappings, which will be needed later. The first is an extension theorem of Kirszbraun (see [34], or [22, Thm. 2.10.43], for example):

Theorem 1.9 (Kirszbraun)**.** *Let $f\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^n$ be an L-Lipschitz mapping. There there is an L-Lipschitz mapping $F\colon \mathbb{R}^d \to \mathbb{R}^n$ such that $F|_A = f$.*

The second one is a differentiability[7] theorem by Rademacher (see, e.g. [41, Thm. 7.3]).

[6]In fact, for any s it cannot increase the Hausdorff s-dimensional measure of any set by more than L^s; see Mattila [41, Thm. 7.5].

[7]By differentiability we always mean the strong (also called Fréchet) differentiability.

Theorem 1.10 (Rademacher)**.** *Let $f\colon \mathbb{R}^d \to \mathbb{R}^n$ be Lipschitz. Then f is differentiable a.e. in $\mathbb{R}^d$.*

Consequently, $\mathrm{Jac}(f)$ is defined a.e., and thus, expressions like $\int_A \mathrm{Jac}(f)\,\mathrm{d}\mathcal{L}$ are well-defined whenever f is Lipschitz and A is measurable.

It is worth noting that the change of variables formula takes a special form in case of bilipschitz mappings (see, e.g., Fremlin [25, Cor. 263F])

Theorem 1.11 (Change of variables for bilipschitz mappings)**.** *Let $A \subset \mathbb{R}^d$ be measurable and $f\colon A \to \mathbb{R}^d$ be bilipschitz. Then the set $f(A)$ is measurable. Moreover, given $\rho\colon f(A) \to \mathbb{R}$, we have*

$$\int_A |\mathrm{Jac}(f)| \cdot (\rho \circ f)\ \mathrm{d}\mathcal{L} = \int_{f(A)} \rho\,\mathrm{d}\mathcal{L}$$

if at least one of the two integrals is defined.

In case that $\rho(x) = 1$ for a.e. $x \in f(A)$ in the theorem above, one obtains the *area formula* for bilipschitz mappings:

$$\mathcal{L}(f(A)) = \int_A |\mathrm{Jac}(f)|\ \mathrm{d}\mathcal{L}.$$

Another easy, but important consequence of Theorem 1.11 is the inverse function theorem for the Jacobians of bilipschitz mappings: given $f\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^d$ bilipschitz, we have

$$\left(\mathrm{Jac}\left(f^{-1}\right) \circ f\right)(x) = \frac{1}{\mathrm{Jac}(f)}(x) \qquad \text{for a.e. } x \in A.$$

We will use a special class of Lipschitz mappings that lie strictly between bilipschitz and general Lipschitz mappings. It was introduced by David [13] and we call them *Lipschitz regular*[8]. The whole Section 1.2 is devoted to the study of such mappings.

Definition 1.12 (Lipschitz regular mapping)**.** *We say that a Lipschitz mapping $f\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^n$ is* Lipschitz regular *if there is $C > 0$ such that for every $y \in \mathbb{R}^n$ and every $r > 0$ the set $f^{-1}(B(y,r))$ can be covered by at most C open balls of radius Cr. The smallest such C is referred to as the* regularity constant *of f and denoted by* $\mathrm{Reg}(f)$*. We use the term (C,L)*-regular mapping *to denote a Lipschitz regular mapping f with Lipschitz constant at most L and* $\mathrm{Reg}(f) \le C$.

We have already mentioned in (1.3) that an L-Lipschitz mapping $f\colon \mathbb{R}^d \to \mathbb{R}^d$ cannot increase the Lebesgue measure of any set more than by a factor of L^d. Considering the dimension fixed, it is just a constant blow-up. The virtue of the Lipschitz regular mappings is, in contrast to the general Lipschitz mappings, that they also cannot compress the measure too much. In fact, they are characterised by this property:

[8] They are called just *regular* mappings in the literature, but since the word 'regular' is heavily overused in mathematics, we choose to expand the name to make it less ambiguous.

Lemma 1.21 (David and Semmes [14]). *Let $f\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^d$ be Lipschitz. Then f is Lipschitz regular if and only if there is $C > 0$ such that*

$$\mathcal{L}(f^{-1}(B(y,r))) \leq Cr^d$$

for every $y \in \mathbb{R}^d$ and every $r > 0$.

Let $\rho\colon I^d \to [0,\infty)$ be measurable. We say that ρ is *bilipschitz realisable*, if there is a bilipschitz mapping $f\colon I^d \to \mathbb{R}^d$ satisfying

$$f_\sharp \rho \mathcal{L} = \mathcal{L}|_{f(I^d)}, \tag{1.4}$$

otherwise we call ρ *bilipschitz non-realisable*. Similarly, we say that ρ is *Lipschitz regular realisable*, if there is a Lipschitz regular solution $f\colon I^d \to \mathbb{R}^d$ to equation (1.4). Otherwise, we call ρ *Lipschitz regular non-realisable.*

Note that for a bilipschitz mapping $f\colon A \to \mathbb{R}^d$ we can rewrite equation (1.4) using the change of variables formula from Theorem 1.11. Given $\rho\colon I^d \to [0,\infty)$, we infer that

$$f_\sharp \rho \mathcal{L}(B) = \int_{f^{-1}(B)} \rho \,\mathrm{d}\mathcal{L} = \int_B \left(\rho \circ f^{-1}\right) \cdot \left|\mathrm{Jac}(f^{-1})\right| \,\mathrm{d}\mathcal{L}$$

for any $B \subseteq f(A)$ measurable. If we use now the inverse function theorem for the Jacobians of bilipschitz mappings, which was stated above, we infer that

$$f_\sharp \rho \mathcal{L}(B) = \int_B \frac{\rho}{|\mathrm{Jac}(f)|} \circ f^{-1} \,\mathrm{d}\mathcal{L}.$$

It follows that $f_\sharp \rho \mathcal{L} = \mathcal{L}|_{f(I^d)}$ if and only if $\rho(x) = |\mathrm{Jac}(f)(x)|$ for a.e. $x \in A$. This explains that equation (1.4) for Lipschitz regular mappings is indeed a generalisation of the prescribed Jacobian equation for bilipschitz mappings stated in equation (1.1).

The Baire category theorem and porosity

Let (X,d) be a metric space. A subset $A \subseteq X$ is called *nowhere dense* in X if $\operatorname{int} \overline{A} = \emptyset$. In other words, for every $x \in X$ and every $r > 0$ there are $x' \in B(x,r)$ and $r' > 0$ such that $B(x',r') \subset B(x,r) \setminus A$. A subset $A \subseteq X$ is called *meagre*[9] if it can be expressed as a countable union of nowhere dense sets.

Nowhere dense and meagre subsets provide a notion of negligible sets in several classes of metric spaces. The class important to the present work consists of all *complete* metric spaces. Their property expressed in the following theorem was proven by Baire [1]:

Theorem 1.13 (The Baire Category Theorem (only a part)). *Let (X,d) be a complete metric space. Let $A \subset X$ be meagre in X. Then $X \setminus A$ is dense in X.*

In particular, the theorem asserts that a complete metric space cannot be expressed as a countable union of nowhere dense sets. The Baire Category Theorem is an important tool in functional analysis, but it has found applications in other

[9] Historically, meagre sets were called 'sets of first category'.

fields of mathematics; a concise treatment of the history of the theorem and its many applications can be found, e.g., in Jones [31].

We add definitions of *porosity* and *σ-porosity* in accordance with Zajíček [49, Def. 2.1], where they are referred to as 'lower porosity' and 'lower σ-porosity', respectively.

Definition 1.14. *Let (X, d) be a metric space.*

1. *A set $P \subseteq X$ is called* porous at a point $x \in X$ *if there exist $\varepsilon_0 > 0$ and $\alpha \in (0, 1)$ such that for every $\varepsilon \in (0, \varepsilon_0)$ there exists $y \in X$ satisfying*

$$d(y, x) \leq \varepsilon \quad \textit{and} \quad B(y, \alpha\varepsilon) \subset B(x, \varepsilon) \setminus P.$$

2. *A set $P \subseteq X$ is called* porous *if P is porous at every point $x \in P$.*

3. *A set $E \subseteq X$ is called* σ-porous *if it may be expressed as a countable union of porous subsets of X.*

Every porous set is nowhere dense, but the opposite implication does not hold. Similarly, in Euclidean spaces every porous set has Lebesgue measure zero, but not vice versa. In fact, Zajíček [48] has shown that there is a set in $\mathbb{R}^d$ that is not σ-porous, but is meagre and of Lebesgue measure zero at the same time. For a survey on porous and σ-porous sets we refer the reader to Zajíček [49].

Due to its relevance in Section 1.4, we point out that porosity of a set $P \subseteq X$ is a weaker condition than requiring P to be porous at all points $x \in X$ (not just at points $x \in P$). For example, the set $\left\{\frac{1}{n} \colon n \in \mathbb{Z} \setminus \{0\}\right\}$ is porous in $\mathbb{R}$, but is not porous at the point 0. However, the corresponding notions of σ-porosity coincide for both versions of porosity discussed here (see Zajíček [49, Prop. 2.5]).

1.1 A discrete and a continuous question

The aim in the present section is to prove Theorem 1.6, which provides a connection between Question 1.4 and its continuous analogue stated in Question 1.5.

We start by providing a reformulation of Question 1.4 that fits better our tools.

Question 1.15. *Is there for every $r > 0$ a constant $L = L(r) > 0$ such that for every $n \in \mathbb{N}$ and every r-separated set $S \subset \mathbb{R}^d$ such that $|S| = n^d$ there is an L-Lipschitz bijection $f\colon S \to \{1, \ldots, n\}^d$?*

Proof of equivalence of Questions 1.4 and 1.15. It is immediate that a negative answer to Question 1.4 provides a negative answer to Question 1.15. Thus, we focus on the opposite direction.

Let $r > 0$ be such that there is no $L(r) > 0$ as in Question 1.15 in dimension $d \in \mathbb{N}$, $d \geq 2$. Let $n \in \mathbb{N}$ and $S \subset \mathbb{R}^d$ be an r-separated set of cardinality n^d. We consider a linear mapping $h\colon \mathbb{R}^d \to \mathbb{R}^d$ defined as $h(x) := \frac{d}{r}x$. Let us write $S' := h(S)$ for a copy of S scaled by the factor $\frac{d}{r}$. The set S' is d-separated.

Next, we define a mapping $z\colon S' \to \mathbb{Z}^d$ by choosing $z(x)$ to be the point $x' \in \mathbb{Z}^d$ that minimises $\|x - x'\|_2$. If there is more than one such point, we choose one of them arbitrarily. Since S' is d-separated, for every $x, y \in S'$ the points $z(x)$ and $z(y)$ are distinct whenever $x \neq y$. We form a set $S'' := z(S') \subset \mathbb{Z}^d$.

Now we assume, for contradiction, that there is $L > 0$ as in Question 1.4. Therefore, there is an L-Lipschitz bijection $f\colon S'' \to [n]^d$. The mapping $f \circ z$ is clearly a bijection. We verify that it is also Lipschitz:

$$\begin{aligned}\|f \circ z(x) - f \circ z(y)\|_2 &\leq L \cdot \|z(x) - z(y)\|_2 \\ &\leq L \cdot (\|x - y\|_2 + \sqrt{d}) \leq L\Big(1 + \sqrt{d}\Big)\,\|x - y\|_2\end{aligned}$$

whenever $x, y \in S'$ are two distinct points. Consequently, the mapping $f \circ z \circ h$ defines an $\frac{Ld\left(1+\sqrt{d}\right)}{r}$-Lipschitz bijection $S \to [n]^d$. Since the last Lipschitz constant does not depend on the original choice of r-separated set S, we get a contradiction. $\square$

The core of the argument proving Theorem 1.6, which is a variant of a discretisation procedure of Burago and Kleiner [8], is contained in the following theorem. It also provides more details on the relation between Questions 1.4 and 1.5.

Theorem 1.16. *Assume that the answer to Question 1.15 is positive. Then for every measurable function $\rho\colon I^d \to [0, \infty)$ such that $\sup \rho < \infty$ and $\rho(x) > 0$ a.e. there is a surjective Lipschitz mapping $f\colon I^d \to \left[0, \sqrt[d]{⨍ \rho}\right]^d$ verifying the equation*

$$f_{\#}\rho\mathcal{L} = \mathcal{L}|_{f(I^d)}. \tag{1.5}$$

Recall that equation (1.5) is equivalent to equation (1.2) from Question 1.5, as follows from the definition of pushforward measure. Theorem 1.6 is an easy consequence of Theorem 1.16, as is shown in the next corollary:

Corollary 1.17. *Let $\rho\colon I^d \to (0, \infty)$ be a measurable function such that $0 < \inf \rho \leq \sup \rho < \infty$. Then any Lipschitz solution $f\colon I^d \to \mathbb{R}^d$ to equation* (1.5) *is also Lipschitz regular.*

Proof. Since we assume that $f_\sharp \rho\mathcal{L} = \mathcal{L}|_{f(I^d)}$, we get for every measurable set $A \subseteq f(I^d)$ that

$$\mathcal{L}(A) = f_\sharp \rho\mathcal{L}(A) = \int_{f^{-1}(A)} \rho \,\mathrm{d}\mathcal{L} \geq \mathcal{L}(f^{-1}(A)) \inf \rho.$$

This implies that $\mathcal{L}(f^{-1}(A)) \leq \frac{\mathcal{L}(A)}{\inf \rho}$. Applying Lemma 1.21 we conclude that f is Lipschitz regular. $\square$

Before we prove Theorem 1.16, let us add a convenient observation, which asserts that it is sufficient to prove Theorem 1.16 only for functions ρ having average value one. Note that the assumptions of Theorem 1.16 imply that the functions ρ considered there have a positive average value.

Observation 1.18. *Let $\rho\colon I^d \to [0, \infty)$ be a measurable function. Then the equation $f_\sharp(\xi\mathcal{L}) = \mathcal{L}|_{f(I^d)}$ admits a Lipschitz solution $f\colon I^d \to \mathbb{R}^d$ for $\xi = \rho$ if and only if it admits Lipschitz solutions for $\xi = \alpha\rho$ for every $\alpha > 0$. Moreover, this equivalence preserves Lipschitz regularity.*

Proof. We write $f\colon I^d \to \mathbb{R}^d$ for a Lipschitz mapping satisfying $f_\sharp(\rho\mathcal{L}) = \mathcal{L}|_{f(I^d)}$. Let us consider a mapping $\phi_\alpha(x) := \sqrt[d]{\alpha} \cdot x$ and observe that $\phi_\alpha \circ f$ is the sought after solution:

$$\begin{aligned}(\phi_\alpha \circ f)_\sharp(\alpha\rho\mathcal{L})(A) &= \alpha \int_{f^{-1}\circ\phi_\alpha^{-1}(A)} \rho \,\mathrm{d}\mathcal{L} = \alpha f_\sharp(\rho\mathcal{L})\Big(\phi_\alpha^{-1}(A)\Big) = \alpha\mathcal{L}|_{f(I^d)}\Big(\phi_\alpha^{-1}(A)\Big) \\ &= \alpha \int_A \left|\operatorname{Jac} \phi_\alpha^{-1}\right| \,\mathrm{d}\mathcal{L}|_{\phi_\alpha\circ f(I^d)} = \mathcal{L}|_{\phi_\alpha\circ f(I^d)}(A)\end{aligned}$$

for any measurable set $A \subseteq (\phi_\alpha \circ f)(I^d)$. The penultimate equality holds by the change of variables (see Theorem 1.11).

Since $\phi_\alpha \circ f$ is only a rescaled version of f, it is clear that f is Lipschitz regular if and only if $\phi_\alpha \circ f$ is. $\square$

The modified Burago–Kleiner discretisation

As we already mentioned, we use a modification of the discretisation procedure of Burago and Kleiner [8] in the presented proof of Theorem 1.16, which they used to encode a bilipschitz non-realisable density ρ into a separated net S in $\mathbb{R}^d$ that cannot be bijectively mapped onto $\mathbb{Z}^d$ in a bilipschitz way. We will use their procedure, with a small technical modification, to encode a given bounded, measurable function $\rho\colon I^d \to [0, \infty)$ into a sequence of separated sets S_i in $\mathbb{R}^d$ such that each S_i has cardinality n_i^d for some $n_i \in \mathbb{N}$.

Burago and Kleiner [8] showed that a bilipschitz bijection $S \to \mathbb{Z}^d$ would yield a bilipschitz solution to equation (1.5). We will show that if there are L-Lipschitz bijections $f_i\colon S_i \to [n_i]^d$ for some $L > 0$ and every $i \in \mathbb{N}$, then there is also a Lipschitz solution to the same equation (1.5). This part of the proof is different than that of Burago and Kleiner, although it follows their overall idea[10].

[10] In fact, in their article Burago and Kleiner [8, Section 2] do not provide the full details for that part of their argument. However, the author was able to verify their result. The arguments presented here are not sufficient in their setting, since in the case we are dealing with the image $f_i(S_i)$ is much nicer.

Let us first describe the main ideas of the modified construction informally. Each set S_i in the sequence of sets 'discretising' a given bounded, measurable function $\rho\colon I^d \to [0,\infty)$ represents 'a picture' of ρ taken with a resolution increasing with i. More precisely, to define S_i we first blow up the domain of ρ by some factor l_i and then subdivide it into m_i^d cubes of the same size. Let T be one of these cubes. The set $S_i \cap T$ will be then formed by regularly spaced points inside T of number proportional to the average value of the blow up of ρ by the factor l_i over T. If we choose l_i and m_i so that $\frac{l_i}{m_i}$ goes to infinity with i, the set S_i will capture more and more precisely variations of ρ on small scales. This idea is illustrated in Figure 1.1.

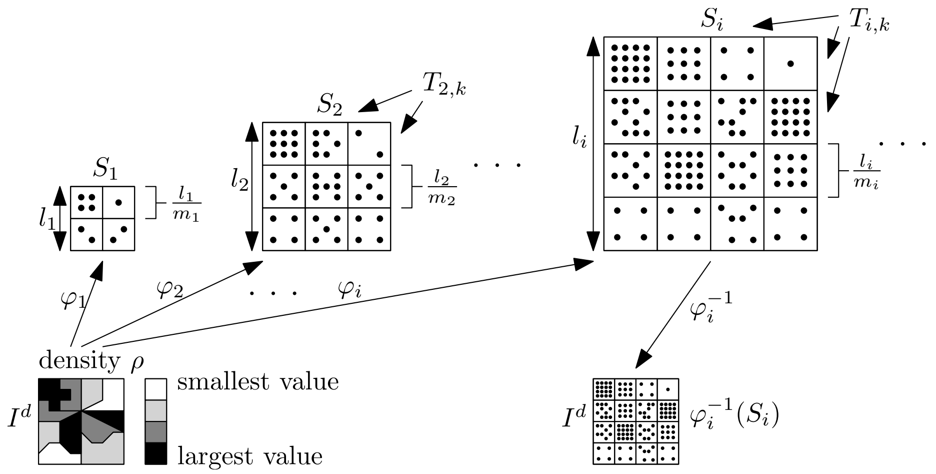

Figure 1.1: An illustration of the construction that encodes a given density ρ into a sequence of separated sets $(S_i)_{i=1}^{\infty}$. After rescaling by φ_i^{-1}, the position and the number of points inside S_i approximate ρ, with a precision increasing with i.

Since we want to use the sets S_i in Question 1.15, we need to make sure that the total number of points forming each S_i is a d-th power of some natural number. That's the technical difference in comparison to the original construction of Burago and Kleiner [8]. Now, we will write everything formally.

We assume that we are given a bounded measurable function $\rho\colon I^d \to [0,\infty)$ with positive average value. We choose two sequences $(l_i)_{i\in\mathbb{N}} \subset \mathbb{R}^+$ and $(m_i)_{i\in\mathbb{N}} \subset \mathbb{N}$. We put several conditions on them, which we describe a bit later. First, let us introduce some notation.

We write $\varphi_i(x) := l_i \cdot x$ for a blow up by factor l_i. It is clear that $\varphi_i(I^d) = [0,l_i]^d$. Each cube $[0,l_i]^d$ naturally decomposes into m_i^d cubes of side $\frac{l_i}{m_i}$; we denote them by $(T_{i,k})_{k=1}^{m_i^d}$. We choose

$$n_{i,k} \in \left\{ \left\lfloor \int_{T_{i,k}} \rho\circ\varphi_i^{-1}\,\mathrm{d}\mathcal{L} \right\rfloor, \left\lfloor \int_{T_{i,k}} \rho\circ\varphi_i^{-1}\,\mathrm{d}\mathcal{L} \right\rfloor + 1 \right\}.$$

This number will stand for $|S_i \cap T_{i,k}|$. The possibility to choose $n_{i,k}$ among two different values will allow us to ensure that $|S_i| = n_i^d$ for some $n_i \in \mathbb{N}$.

The change of variables formula (see Theorem 1.11) implies

$$\int_{T_{i,k}} \rho\circ\varphi_i^{-1}\,\mathrm{d}\mathcal{L} = \int_{\varphi_i^{-1}(T_{i,k})} \rho\,|\mathrm{Jac}(\varphi_i)|\,\mathrm{d}\mathcal{L} = l_i^d \int_{\varphi_i^{-1}(T_{i,k})} \rho\,\mathrm{d}\mathcal{L}. \tag{1.6}$$

Using the upper bound on ρ we infer that

$$n_{i,k} \leq 1 + \left\lfloor \frac{\sup \rho \cdot l_i^d}{m_i^d} \right\rfloor \tag{1.7}$$

Now we can state the required conditions on l_i, m_i and $n_{i,k}$:

1. $l_i \to \infty$, $m_i \to \infty$ and $\frac{l_i}{m_i} \to \infty$ as $i \to \infty$.
2. for every $i \in \mathbb{N}$ we choose each $n_{i,k}$ from the two possibilities so that there is $n_i \in \mathbb{N}$ such that $n_i^d = \sum_{k=1}^{m_i^d} n_{i,k}$.

We show that these conditions can be satisfied at once. We set $l_i := m_i^{1+p}$ for a suitable $p > 0$ and choose $(m_i)_{i=1}^{\infty} \subset \mathbb{N}$ as an increasing sequence. This will satisfy the first condition.

In order to satisfy the second condition, it is sufficient to make sure that the interval

$$\left[\sum_{k=1}^{m_i^d} \left\lfloor \int_{T_{i,k}} \rho \circ \varphi_i^{-1} \, \mathrm{d}\mathcal{L} \right\rfloor, m_i^d + \sum_{k=1}^{m_i^d} \left\lfloor \int_{T_{i,k}} \rho \circ \varphi_i^{-1} \, \mathrm{d}\mathcal{L} \right\rfloor \right]$$

contains a d-th power of a natural number. If we denote by a_i the largest integer such that $a_i^d < \sum_{k=1}^{m_i^d} \left\lfloor \int_{T_{i,k}} \rho \circ \varphi_i^{-1} \, \mathrm{d}\mathcal{L} \right\rfloor$, we need that $(a_i + 1)^d \leq m_i^d + \sum_{k=1}^{m_i^d} \left\lfloor \int_{T_{i,k}} \rho \circ \varphi_i^{-1} \, \mathrm{d}\mathcal{L} \right\rfloor$. Since $(a_i + 1)^d - a_i^d \leq C(d) a_i^{d-1}$, where $C(d)$ is a constant depending only on the dimension d, it is sufficient to choose m_i so that $m_i^d > C(d) a_i^{d-1}$. From equation (1.6) we get that

$$a_i^d < \sup \rho \cdot l_i^d = \sup \rho \cdot m_i^{(1+p)d};$$

thus we see that m_i satisfies $m_i^d > C(d) a_i^{d-1}$ provided we choose $p < \frac{1}{d-1}$ and m_1 sufficiently large[11].

After setting up the parameters l_i and m_i properly, we can construct the separated sets S_i. We first form sets $S_{i,k}$ by placing $n_{i,k}$ distinct points inside each $T_{i,k}$ and then set $S_i := \bigcup_{k=1}^{m_i^d} S_{i,k}$. But instead of providing an explicit formula for $S_{i,k}$, it will be enough to consider any sufficiently separated set of $n_{i,k}$ points inside $T_{i,k}$ and argue that the separation constant can be chosen independently of i and k.

Since each $T_{i,k}$ has a side of length $\frac{l_i}{m_i}$, given $r > 0$ satisfying

$$\frac{l_i}{m_i \lceil \sqrt[d]{n_{i,k}} \rceil} \geq r, \tag{1.8}$$

we may define $S_{i,k}$ as any r-separated set of $n_{i,k}$ points inside $T_{i,k}$ that, in addition, satisfies $\operatorname{dist}(S_{i,k}, \partial T_{i,k}) \geq r/2$; an example is depicted in Figure 1.2. The last condition ensures that the set S_i is r-separated as well. Note that in equation (1.8) we assumed that $n_{i,k}$ is at least one. If it is zero, we simply place zero points inside $T_{i,k}$.

[11] We note that it is necessary to put an additional mild assumption on the value of m_1 in order for the numbers a_i to be well defined already starting with $i = 1$; namely, we assume that m_1 is big enough for $\sum_{k=1}^{m_1^d} \left\lfloor \int_{T_{1,k}} \rho \circ \varphi_1^{-1} \, \mathrm{d}\mathcal{L} \right\rfloor$ to be larger than one.

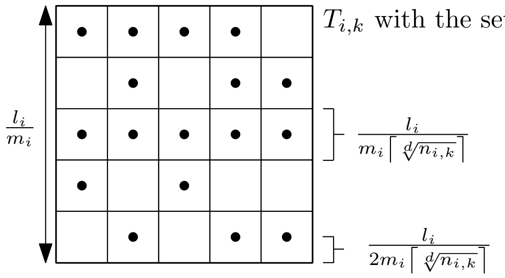

Figure 1.2: The construction of $S_{i,k}$ inside $T_{i,k}$.

It remains to verify the existence of the separation constant $r > 0$ satisfying the inequality (1.8) for all i, k. Using the inequality (1.7) and standard estimates we obtain that

$$\lceil \sqrt[d]{n_{i,k}}\rceil^d \leq 2^d n_{i,k} \leq 2^d\left(\frac{\sup\rho\cdot l_i^d}{m_i^d}+1\right)$$

for $n_{i,k} \geq 1$ (clearly, cubes $T_{i,k}$ containing zero points do not decrease the separation of the set S_i). With an additional mild assumption on m_1, namely that $\frac{\sup\rho\cdot l_1^d}{m_1^d} \geq 1$, we infer that

$$2^d\left(\frac{\sup\rho\cdot l_i^d}{m_i^d}+1\right) \leq 2^{2d}\frac{\sup\rho\cdot l_i^d}{m_i^d}.$$

This in turn provides us with the bound $\lceil \sqrt[d]{n_{i,k}}\rceil \leq \frac{4\sqrt[d]{\sup\rho}\cdot l_i}{m_i}$. Substituting this bound into the inequality (1.8), we see that we may take $r := \frac{1}{4\sqrt[d]{\sup\rho}}$. This finishes the description of the construction of the separated sets S_i.

The proof of Theorem 1.16

We are now going to employ the construction presented in the previous part to prove Theorem 1.16. However, we need to present two auxiliary lemmas on weak convergence of measures first, which are probably a part of a common knowledge in measure theory, but the author was unable to find a proper reference.

Lemma 1.19. *Let ν and $(\nu_n)_{n=1}^\infty$ be finite Borel measures with support in a compact set $K \subset \mathbb{R}^d$. Moreover, assume that there is, for each $n \in \mathbb{N}$, a finite collection $\mathcal{Q}_n$ of Borel subsets of K that satisfy the following:*

1. $\nu\left(K \setminus \bigcup \mathcal{Q}_n\right) = 0$ *and* $\nu_n\left(K \setminus \bigcup \mathcal{Q}_n\right) = 0$.

2. $\sum_{Q\in\mathcal{Q}_n} \nu(Q) = \nu(K)$ *and* $\sum_{Q\in\mathcal{Q}_n} \nu_n(Q) = \nu_n(K)$.

3. $\lim_{n\to\infty} \max_{Q\in\mathcal{Q}_n} \operatorname{diam}(Q) = 0$ *and* $\max_{Q\in\mathcal{Q}_n} |\nu_n(Q) - \nu(Q)| \in o\left(\frac{1}{|\mathcal{Q}_n|}\right)$.

Then ν_n converges weakly to ν.

Proof. Fix any $\psi \in C(K)$ and $\varepsilon > 0$. For every $Q \in \mathcal{Q}_n$ we choose $z_Q \in Q$ arbitrarily. By the uniform continuity of ψ, there is $N \in \mathbb{N}$ such that for every $n \geq N$, every $Q \in \mathcal{Q}_n$ and every $x \in Q$ we have $|\psi(z_Q) - \psi(x)| \leq \varepsilon$. Moreover, we require that for every $n \geq N$ it holds that

$$\max_{Q \in \mathcal{Q}_n} |\nu_n(Q) - \nu(Q)| \leq \frac{\varepsilon}{|\mathcal{Q}_n|}.$$

We write

$$\begin{aligned}
\int_Q \psi \, \mathrm{d}\nu_n &\leq \int_Q \psi(z_Q) + \varepsilon \, \mathrm{d}\nu_n = (\psi(z_Q) + \varepsilon)\nu_n(Q) \leq (\psi(z_Q) + \varepsilon)\left(\nu(Q) + \frac{\varepsilon}{|\mathcal{Q}_n|}\right) \\
&= \int_Q (\psi(z_Q) + \varepsilon) \, \mathrm{d}\nu + (\psi(z_Q) + \varepsilon)\frac{\varepsilon}{|\mathcal{Q}_n|} \\
&\leq \int_Q (\psi + 2\varepsilon) \, \mathrm{d}\nu + (\psi(z_Q) + \varepsilon)\frac{\varepsilon}{|\mathcal{Q}_n|} \\
&\leq \int_Q \psi \, \mathrm{d}\nu + 2\varepsilon\nu(Q) + (\psi(z_Q) + \varepsilon)\frac{\varepsilon}{|\mathcal{Q}_n|}.
\end{aligned}$$

Symmetrically, we derive the lower bound

$$\int_Q \psi \, \mathrm{d}\nu_n \geq \int_Q \psi \, \mathrm{d}\nu - 2\varepsilon\nu(Q) - (\psi(z_Q) - \varepsilon)\frac{\varepsilon}{|\mathcal{Q}_n|}.$$

Summing over all elements of $\mathcal{Q}_n$ and using assumptions 1 and 2, ensuring that elements of $\mathcal{Q}_n$ are almost disjoint and cover almost all of K with respect to both measures ν and ν_n, we deduce

$$\left| \int_K \psi \, \mathrm{d}\nu_n - \int_K \psi \, \mathrm{d}\nu \right| \leq 2\varepsilon\nu(K) + \varepsilon(\max |\psi| + \varepsilon).$$

Since ψ is continuous on a compact K, it is also bounded and the right hand side of the last inequality tends to zero with ε. $\square$

Lemma 1.20. *Let K be a compact set in $\mathbb{R}^d$ and $(\nu_n)_{n \in \mathbb{N}}$ be a sequence of finite Borel measures on K converging weakly to a finite Borel measure ν. Let $(h_n)_{n \in \mathbb{N}}$, $h_n \colon K \to \mathbb{R}^m$, be a sequence of continuous mappings converging uniformly to a mapping h. Then $(h_n)_\sharp(\nu_n)$ converges weakly to $h_\sharp(\nu)$.*

Proof. We take any $\psi \in C(\mathbb{R}^m)$ with compact support. Using the abstract version of change of variables (see Theorem 1.7), we can bound

$$\begin{aligned}
&\left| \int_{h_n(K)} \psi \, \mathrm{d}(h_n)_\sharp(\nu_n) - \int_{h(K)} \psi \, \mathrm{d}h_\sharp(\nu) \right| = \left| \int_K \psi \circ h_n \, \mathrm{d}\nu_n - \int_K \psi \circ h \, \mathrm{d}\nu \right| \\
&\leq \int_K |\psi \circ h_n - \psi \circ h| \, \mathrm{d}\nu_n + \left| \int_K \psi \circ h \, \mathrm{d}\nu_n - \int_K \psi \circ h \, \mathrm{d}\nu \right|.
\end{aligned}$$

As $n \to \infty$ the first term in the final sum goes to zero since $\psi \circ h_n$ converges uniformly to $\psi \circ h$. Moreover, the second term converges to zero as well, because ν_n converges weakly to ν. $\square$

We are ready to prove Theorem 1.16.

Proof of Theorem 1.16. We assume that we are given a bounded measurable function $\rho\colon I^d \to [0,\infty)$ that is positive a.e. As noted before, it implies that its average value is positive as well. Using Observation 1.18, we may assume that $\fint_{I^d} \rho = 1$. We use the construction described in the previous part on ρ and get sequences of r-separated sets S_i and $S_{i,k}$ together with the parameters $l_i, m_i, n_i, n_{i,k}$, mappings φ_i and cubes $T_{i,k}$.

Assuming the positive answer to Question 1.15, we get $L > 0$ and a sequence of L-Lipschitz bijections $f_i\colon S_i \to [n_i]^d$. We pull each f_i back to I^d in the following way. We write X_i for the set $\varphi_i^{-1}(S_i)$ and define a mapping $g_i\colon X_i \to \mathbb{R}^d$ as $g_i(x) := \frac{1}{n_i} \cdot f_i \circ \varphi_i(x)$. It is not hard to see that the mappings g_i are uniformly Lipschitz:

$$\begin{aligned}\|g_i(x) - g_i(y)\|_2 &\le \frac{1}{n_i} \|f_i \circ \varphi_i(x) - f_i \circ \varphi_i(y)\|_2 \\ &\le \frac{L}{n_i} \|\varphi_i(x) - \varphi_i(y)\|_2 = L\frac{l_i}{n_i} \|x - y\|_2 ,\end{aligned}$$

for every $x, y \in X_i$. Thus, we need to examine the behaviour of the sequence $\left(\frac{l_i}{n_i}\right)_{i=1}^{\infty}$.

From the definition of $n_{i,k}$ we have the following bounds on $n_i^d = \sum_{k=1}^{m_i^d} n_{i,k}$:

$$\int_{\varphi_i(I^d)} \rho \circ \varphi_i^{-1} \, \mathrm{d}\mathcal{L} + m_i^d \ge n_i^d \ge \int_{\varphi_i(I^d)} \rho \circ \varphi_i^{-1} \, \mathrm{d}\mathcal{L} - m_i^d.$$

Using the identity $\fint_{I^d} \rho = 1$ and the fact that $\mathrm{Jac}(\varphi_i) = l_i^d$ we obtain that

$$l_i^d + m_i^d \ge n_i^d \ge l_i^d - m_i^d.$$

Since $\frac{m_i}{l_i} \to 0$, the sequence $\left(\frac{n_i}{l_i}\right)_{i \in \mathbb{N}}$ is bounded and converges to 1, and hence, $\frac{l_i}{n_i} \to 1$ as well. Consequently, for any $L' > L$ we can find $i_0 \in \mathbb{N}$ such that for every $i \ge i_0$ the mappings g_i are L'-Lipschitz. By trimming off the initial segment of the sequence $(g_i)_{i \in \mathbb{N}}$ up to i_0, we can assume that all mappings g_i are L'-Lipschitz for any chosen $L' > L$.

We extend each g_i by Kirszbraun's extension theorem (see Theorem 1.9 and the references therein) to a mapping $\bar{g}_i\colon I^d \to \mathbb{R}^d$ such that $\mathrm{Lip}(\bar{g}_i) = \mathrm{Lip}(g_i)$. By the Arzelà–Ascoli theorem (see, e.g., Dunford and Schwartz [19, Thm. IV.6.7]) we know that the sequence $(\bar{g}_i)_{i=1}^{\infty}$ subconverges to a limit f, which is also L'-Lipschitz. In fact, it is L-Lipschitz, as follows from the previous discussion. By passing to a convergent subsequence, we may assume that $\bar{g}_i \rightrightarrows f$.

For $i \ge 1$ we define a measure μ_i on I^d by

$$\mu_i(A) = \frac{1}{n_i^d} |A \cap X_i| , \qquad A \subseteq I^d.$$

In order to show that $f_\sharp(\rho\mathcal{L}) = \mathcal{L}|_{I^d}$ we first prove that μ_i converges weakly to $\rho\mathcal{L}$ on I^d. Moreover, this will be shown to imply that $(\bar{g}_i)_\sharp(\mu_i)$ converges weakly to $f_\sharp(\rho\mathcal{L})$. Finally, we prove that $(\bar{g}_i)_\sharp(\mu_i)$ also converges weakly to $\mathcal{L}|_{I^d}$, and hence, $f_\sharp(\rho\mathcal{L})$ and $\mathcal{L}|_{I^d}$ must be the same by the uniqueness of weak limits.

Claim 1.16.1. *The sequence of measures $(\mu_i)_{i \in \mathbb{N}}$ converges weakly to $\rho\mathcal{L}$.*

Proof. We just need to verify that the measures μ_i and $\rho\mathcal{L}$ satisfy the assumptions of Lemma 1.19. The only non-trivial assumption in this case is the existence of the collection $\mathcal{Q}_i$. We take $\mathcal{Q}_i := \left\{\varphi_i^{-1}(T_{i,k})\colon k \in [m_i^d]\right\}$. By construction, the sets $\varphi_i^{-1}(T_{i,k})$ form a decomposition of I^d into m_i^d cubes of side $1/m_i$. Since m_i goes to infinity with i, the diameter of $\varphi_i^{-1}(T_{i,k})$ goes to zero.

Clearly, the overlap of any two of these cubes has measure zero, and thus, $\rho\mathcal{L}$-measure zero as well. Moreover, since the sets S_i were constructed to have nonzero distance to all sets of the form $\partial T_{i,k}$, the μ_i-measure of the overlap of any two cubes from $\mathcal{Q}_i$ must be zero, too.

It remains to check that $\max_{Q\in\mathcal{Q}_i} |\mu_i(Q) - \rho\mathcal{L}(Q)| \in o\Big(\frac{1}{m_i^d}\Big)$. To see this, recall that μ_i is supported on $\varphi_i^{-1}(T_{i,k})$ by the set $\varphi_i^{-1}(T_{i,k}) \cap X_i$ consisting of $n_{i,k}$ points. For any $i \in \mathbb{N}, k \in [m_i^d]$ we can write

$$\begin{aligned}\mu_i\left(\varphi_i^{-1}(T_{i,k})\right) =& \frac{n_{i,k}}{n_i^d} \le \frac{1}{n_i^d}\left(\left\lfloor \int_{T_{i,k}} \rho\circ\varphi_i^{-1}\,\mathrm{d}\mathcal{L}\right\rfloor + 1\right)\\ \le& \frac{l_i^d}{n_i^d}\int_{\varphi_i^{-1}(T_{i,k})} \rho\,\mathrm{d}\mathcal{L} + \frac{1}{n_i^d} \le \frac{l_i^d}{n_i^d}\rho\mathcal{L}\left(\varphi_i^{-1}(T_{i,k})\right) + \frac{1}{n_i^d},\end{aligned}$$

and similarly,

$$\mu_i\left(\varphi_i^{-1}(T_{i,k})\right) \ge \frac{1}{n_i^d}\left(\int_{T_{i,k}} \rho\circ\varphi_i^{-1}\,\mathrm{d}\mathcal{L} - 1\right) \ge \frac{l_i^d}{n_i^d}\rho\mathcal{L}\left(\varphi_i^{-1}(T_{i,k})\right) - \frac{1}{n_i^d},$$

Therefore, we can bound $\left|\mu_i\left(\varphi_i^{-1}(T_{i,k})\right) - \rho\mathcal{L}\left(\varphi_i^{-1}(T_{i,k})\right)\right|$ above as

$$\frac{1}{n_i^d} + \left|\frac{l_i^d}{n_i^d}\rho\mathcal{L}\left(\varphi_i^{-1}(T_{i,k}\right) - \rho\mathcal{L}\left(\varphi_i^{-1}(T_{i,k})\right)\right| = \frac{1}{n_i^d} + \rho\mathcal{L}\left(\varphi_i^{-1}(T_{i,k})\right)\left|\frac{l_i^d}{n_i^d} - 1\right|.$$

Recalling that $\rho\mathcal{L}\left(\varphi_i^{-1}(T_{i,k})\right) \le \frac{\sup\rho}{m_i^d}$, $\frac{l_i}{m_i} \to \infty$ and $\frac{l_i}{n_i} \to 1$ as $i \to \infty$, this proves that

$$\left|\mu_i\left(\varphi_i^{-1}(T_{i,k})\right) - \rho\mathcal{L}\left(\varphi_i^{-1}(T_{i,k})\right)\right| \in o\left(\frac{1}{m_i^d}\right).$$

Hence, μ_i and $\rho\mathcal{L}$ satisfy the assumptions of Lemma 1.19. □

Claim 1.16.2. *The sequence of measures $((\bar{g}_i)_\sharp(\mu_i))_{i\in\mathbb{N}}$ converges weakly to the measure $f_\sharp(\rho\mathcal{L})$.*

Proof. We know that $\bar{g}_i \rightrightarrows f$ and that μ_i converge weakly to $\rho\mathcal{L}$ by Claim 1.16.1. Thus, it is sufficient to directly apply Lemma 1.20. □

Claim 1.16.3. *The sequence of measures $((\bar{g}_i)_\sharp(\mu_i))_{i\in\mathbb{N}}$ converges weakly to $\mathcal{L}|_{I^d}$.*

Proof. Note that for every $i \in \mathbb{N}$ the set $f_i(S_i)$ is exactly the set $[n_i]^d$. Therefore, the set $g_i(X_i)$, which is the support of the measure $(\bar{g}_i)_\sharp(\mu_i)$, is precisely the set $\left\{\frac{1}{n_i}, \frac{2}{n_i}, \ldots, \frac{n_i}{n_i}\right\}^d$, that is, a regular grid with n_i^d points inside I^d. The situation is depicted in Figure 1.3.

Since the weight assigned to each point of X_i by μ_i is exactly $\frac{1}{n_i^d}$, it is intuitively clear that $(\bar{g}_i)_\sharp(\mu_i)$ converges weakly to $\mathcal{L}|_{I^d}$. For a formal justification, the conditions of Lemma 1.19 are easily verified for $\nu_n := (\bar{g_n})_\sharp(\mu_n)$, $\nu := \mathcal{L}|_{I^d}$ and

$$\mathcal{Q}_n := \left\{ \prod_{j=1}^{d} \left(\frac{b_j - 1}{n_i}, \frac{b_j}{n_i} \right] : (b_1, \ldots, b_d) \in [n_i]^d \right\}.$$

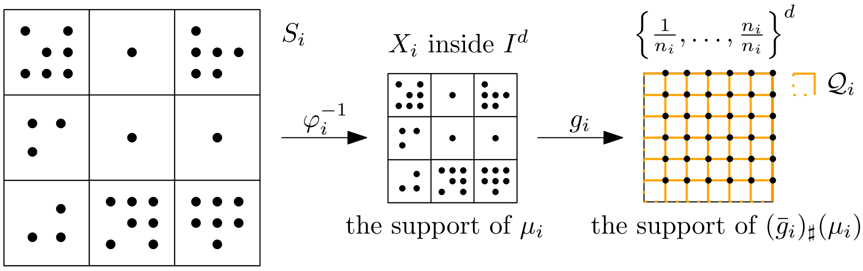

Figure 1.3: Weak convergence of $(\bar{g}_i)_\sharp(\mu_i)$ to $\mathcal{L}|_{I^d}$.

Combining Claim 1.16.2 and the last observation we infer that $f_\sharp(\rho\mathcal{L}) = \mathcal{L}|_{I^d}$ by the uniqueness of weak limits. $\square$

It remains to observe that $f(I^d) = I^d$. The equation $f_\sharp(\rho\mathcal{L}) = \mathcal{L}|_{I^d}$ and $\rho > 0$ a.e. imply that for any open set $A \subseteq I^d$, the preimage $f^{-1}(A)$ has a positive $\mathcal{L}|_{I^d}$-measure, and thus, is not empty. Since $f(I^d)$ is closed, it follows that I^d is contained in $f(I^d)$.

Similarly, for any set $A \subseteq I^d$ of positive measure, the image $f(A)$ has positive $\mathcal{L}|_{I^d}$-measure. Consequently, $f(A) \cap I^d$ is not empty. Since I^d is closed, we deduce that $f(I^d)$ is contained in I^d as well. $\square$

1.2 Lipschitz regular mappings

The main purpose of the present section is to prove Theorem 1.35, which we need to resolve Feige's question. As we noted in the introduction to the present chapter, Theorem 1.35 can be deduced quite quickly from the work of Bonk and Kleiner [5]. However, Dymond, the author, and Kopecká [21] have developed a completely independent proof, which is simpler and shorter than the proof by Bonk and Kleiner [5], though less general; we present it in this section as well. Beyond that, we present a few complementary results on Lipschitz regular mappings which are not related to Feige's question.

We restate the definition of Lipschitz regular mappings:

Definition 1.12 (Lipschitz regular mapping)**.** *We say that a Lipschitz mapping* $f\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^n$ *is* Lipschitz regular *if there is* $C > 0$ *such that for every* $y \in \mathbb{R}^n$ *and every* $r > 0$ *the set* $f^{-1}(B(y,r))$ *can be covered by at most* C *open balls of radius* Cr*. The smallest such* C *is referred to as the* regularity constant *of* f *and denoted by* $\operatorname{Reg}(f)$*. We use the term* (C, L)-regular mapping *to denote a Lipschitz regular mapping* f *with Lipschitz constant at most* L *and* $\operatorname{Reg}(f) \leq C$*.*

Because the definition of Lipschitz regular mappings uses open balls, we set up a convention that every ball is assumed to be open if not stated otherwise.

Lipschitz regular mappings constitute an intermediate class between Lipschitz and bilipschitz mappings. While bilipschitz mappings are sometimes too rigid, Lipschitz mappings can be very degenerate; they can map many points onto a single one or map sets of positive measure onto sets of measure zero. Various classes of mappings lying somewhere in between Lipschitz and bilipschitz have been studied; for instance, in works of Bates et al. [2], Johnson et al. [29], Maleva [38], or of Benyamini and Lindenstrauss [3, Ch. 11]. Lipschitz regular mappings were introduced, for the first time, by David [13]; see the book by David and Semmes [14, Ch. 2] for a further reference.

The definition and many properties of Lipschitz regular mappings can be stated for metric spaces without any additional difficulties in the proofs. However, in the present work we are interested only in the case of Euclidean spaces. Therefore, we have chosen to restrict all definitions and statements to that setting.

All bilipschitz mappings are Lipschitz regular. A classic example of a non-bilipschitz (in fact non-injective) Lipschitz regular mapping is given by a folding mapping of the plane $\mathbb{R}^2$, i.e., take the plane and fold it along the y-axis. This defines a mapping $\mathbb{R}^2 \to \mathbb{R}^2$ which is Lipschitz regular with regularity constant 2.

David and Semmes studied Lipschitz regular mappings in the context of general metric spaces and Euclidean spaces. In the Euclidean space setting, David proved that Lipschitz regular mappings behave somewhat like bilipschitz mappings. More precisely, that inside any ball B in the domain of a Lipschitz regular mapping $f\colon \mathbb{R}^d \to \mathbb{R}^m$ with $m \geq d$, one can find a set E of large measure so that the restriction of f to the set E is bilipschitz; see David [13, Prop. 1, p.95] or David and Semmes [15, Thm 4.1, p.380]. Although the set E is large in measure, we point out that it may have empty interior. There are two natural questions arising from this result:

Question. *Can one find a non-empty, open ball inside* B *on which* f *is bilipschitz and, if yes, can one additionally demand that the set* E *above is open?*

We provide answers to these questions in the case that the dimension of the domain is equal to the dimension of the co-domain. The main result proven in the present section (Theorem 1.30), which is the answer to the first part of the above question, can be derived quickly from a result of Bonk and Kleiner [5, Thm 3.4 and Lem. 4.2], as we already noted before.

The material at the end of the present section appearing under the heading 'Optimality of Theorem 1.30' can be seen as a complement of Theorem 1.30 and theorem 1.35 and is independent of Feige's question (Question 1.4) and of the work of Bonk and Kleiner [5].

Preliminaries

Before we start the exposition of the results we list general properties of Lipschitz regular mappings that will be needed later. For reader's convenience, we repeat here the equivalent characterisation of Lipschitz regular mappings stated in the 'Background and notation' at the beginning:

Lemma 1.21 (David and Semmes [14])**.** *Let $f\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^d$ be Lipschitz. Then f is Lipschitz regular if and only if there is $C > 0$ such that*

$$\mathcal{L}(f^{-1}(B(y,r))) \leq Cr^d$$

for every $y \in \mathbb{R}^d$ and every $r > 0$.

A particularly useful special case of Lemma 1.21 asserts that Lipschitz regular mappings possess Luzin's properties (N) and (N^{-1}), as stated in the next Corollary:

Corollary 1.22. *Let $f\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^d$ be Lipschitz regular. Then*

(N) *for every $E \subset A$ such that $\mathcal{L}(E) = 0$ we have $\mathcal{L}(f(E)) = 0$; and*

(N^{-1}) *for every $F \subset \mathbb{R}^d$ such that $\mathcal{L}(F) = 0$ we have $\mathcal{L}(f^{-1}(F)) = 0$.*

We also add an easy observation, which, however, will prove useful later.

Observation 1.23. *Let $f\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^d$ be Lipschitz regular and $y \in \mathbb{R}^d$. Then we have*

$$\left|f^{-1}(\{y\})\right| \leq \operatorname{Reg}(f).$$

Proof. To the contrary, we assume there are pairwise distinct points $x_1, \ldots, x_k \in f^{-1}(\{y\})$ for some $k > \operatorname{Reg}(f)$. Let us denote by r the minimum distance between x_i and x_j for $1 \leq i < j \leq k$. Then no ball in $\mathbb{R}^d$ of radius $\frac{r}{2}$ can contain more than one of the points $x_1, \ldots, x_k$. Therefore, $f^{-1}(B(y, \frac{r}{2\operatorname{Reg}(f)}))$ cannot be covered with at most $\operatorname{Reg}(f)$ balls of radius $\frac{r}{2}$, a contradiction. □

The proofs in this section rely on two main tools. The first one is differentiability of Lipschitz regular mappings.

Definition 1.24 (Non-critical points)**.** *Let $U \subseteq \mathbb{R}^d$ be open and $f\colon U \to \mathbb{R}^d$ be Lipschitz regular. We define the set of 'non-critical' points of f as*

$$N(f) := U \setminus f^{-1}(\{f(x) : Df(x) \text{ does not exist or does not have full rank}\})$$

For a mapping f as in the definition above, by a variant of Sard's theorem for Lipschitz mappings, which can be found, e.g., in Mattila [41, Thm. 7.6], we know that the set of 'critical values'

$$\{f(x) : Df(x) \text{ does not exist or does not have full rank}\}$$

has zero Lebesgue measure. Therefore, by Corollary 1.22, the set of 'non-critical points' $N(f)$ occupies almost all of U. Notice that for every $x \in N(f)$ we have that $Df(x)$ exists and is invertible[12], and moreover, that $f^{-1}(\{f(x)\}) \subseteq N(f)$.

Occasionally, we will be given an open set U and a Lipschitz regular mapping f defined on $\overline{U}$, the closure of U. Then by $N(f)$ we mean the set $N(f|_U) \subseteq U$. Note that it is then still true that $f(N(f))$ has full measure in $f(U)$—we will use this fact several times.

Another important tool that is used the present section, besides differentiability, is the notion of topological degree. We briefly introduce it here; for its detailed treatment we refer to Deimling [16, Chaps. 1-2] or Fonseca and Gangbo [24].

The degree function

$$\deg\colon \left\{(f,U,y)\colon \begin{matrix} U \subseteq \mathbb{R}^d \text{ open and bounded,} \\ f \in C(\overline{U};\mathbb{R}^d),\ y \in \mathbb{R}^d \setminus f(\partial U) \end{matrix}\right\} \longrightarrow \mathbb{Z}$$

is uniquely determined by the following three properties [16, Thm. 1.1]:

(d1) $\deg(\mathrm{id}, U, y) = 1$ for all $y \in U$.

(d2) (additivity) $\deg(f, U, y) = \deg(f, U_1, y) + \deg(f, U_2, y)$ whenever U_1, U_2 are disjoint open subsets of U such that $y \notin f\left(\overline{U} \setminus (U_1 \cup U_2)\right)$.

(d3) (homotopy invariance) $\deg(H(t), U, y(t)) = \deg(H(0), U, y(0))$ whenever the mappings

$$H\colon [0,1] \to C(\overline{U}, \mathbb{R}^d) \qquad y\colon [0,1] \to \mathbb{R}^d$$

are continuous and $y(t) \notin H(t)(\partial U)$ for all $t \in [0,1]$.

The degree function is defined explicitly in [16, Chap. 2]. We just point out that in the special case where $g \in C^1(\overline{U}, \mathbb{R}^d)$ and for every point $x \in g^{-1}(\{y\})$ the derivative $Dg(x)$ is invertible, then the degree function is given by the expression

$$\deg(g, U, y) = \sum_{x \in g^{-1}(\{y\})} \mathrm{sign}(\mathrm{Jac}(g)(x)), \tag{1.9}$$

(see [16, Def. 2.1]). In particular, we have that $\deg(g, U, y) = 0$ whenever $y \in \mathbb{R}^d \setminus g(\overline{U})$.

We will require some further properties of the degree which follow easily from the properties (d1), (d2) and (d3). All of the statements of the next Proposition are contained in [16, Thm. 3.1].

[12] In fact, for a Lipschitz regular mapping f as above, it is easy to prove that $Df(x)$ is always invertible whenever it exists; if not, then there would be a point x and a direction v such that the distances between x and points of the form $x+tv$ for $t > 0$ small enough would be contracted by f by an arbitrarily large factor, which would eventually contradict the Lipschitz regularity of f.

Proposition 1.25. *Let $U \subseteq \mathbb{R}^d$ be an open, bounded set, $f \in C(\overline{U}, \mathbb{R}^d)$ and $y \in \mathbb{R}^d \setminus f(\partial U)$.*

(i) If y and y' belong to the same connected component of $\mathbb{R}^d \setminus f(\partial U)$ then $\deg(f, U, y) = \deg(f, U, y')$.

(ii) If $y \in \mathbb{R}^d \setminus f(\overline{U})$ then $\deg(f, U, y) = 0$.

(iii) Let $g \in C(\overline{U}, \mathbb{R}^d)$ be a mapping such that $\|f - g\|_\infty < \operatorname{dist}(y, f(\partial U))$. Then $\deg(f, U, y) = \deg(g, U, y)$.

In the next Proposition, we extend formula (1.9) to Lipschitz mappings. The author was not able to find it in the literature, although it is probably known to experts in the field.

Proposition 1.26. *Let $U \subseteq \mathbb{R}^d$ be an open, bounded set, $f\colon \overline{U} \to \mathbb{R}^d$ be a Lipschitz mapping and $y \in \mathbb{R}^d \setminus f(\partial U)$ be such that for every $x \in f^{-1}(\{y\})$ the derivative $Df(x)$ exists and is invertible. Then*

$$\deg(f, U, y) = \sum_{x \in f^{-1}(\{y\})} \operatorname{sign}(\operatorname{Jac}(f)(x)).$$

Proof. If $f^{-1}(\{y\}) = \emptyset$ then $\deg(f, U, y) = 0$, by Proposition 1.25, part (ii), and the formula holds. Thus, we assume that $f^{-1}(\{y\}) \neq \emptyset$. Note that $f^{-1}(\{y\})$ is finite. Otherwise we may find an accumulation point x of $f^{-1}(\{y\})$. Then $f(x) = y$ and there is a sequence $(x_n)_{n=1}^\infty$ in $f^{-1}(\{y\}) \setminus \{x\}$ such that $x_n \to x$ as $n \to \infty$. Since $f(x_n) - f(x) = 0$, by differentiability at x we have $\|Df(x)(x_n - x)\|_2 \in o(\|x_n - x\|_2)$. It follows that $\inf_{y \in S^{d-1}} \|Df(x)(y)\|_2 = 0$, and thus, $Df(x)$ is not invertible.

Let us write $f^{-1}(\{y\}) = \{x_1, x_2, \dots, x_n\}$ where $n \in \mathbb{N}$. Fix $i \in [n]$. Given $\varepsilon > 0$, we may choose $\delta_i > 0$ sufficiently small so that

$$\|f(x) - f(x_i) - Df(x_i)(x - x_i)\|_2 \leq \varepsilon \|x - x_i\|_2$$

for all $x \in \overline{B}(x_i, \delta_i) \subseteq U$. If we define an affine mapping $g\colon \overline{U} \to \mathbb{R}^d$ by $g(x) = f(x_i) + Df(x_i)(x - x_i)$, the inequality above yields $\|f(x) - g(x)\|_2 \leq \varepsilon \|x - x_i\|_2$ for every $x \in \overline{B}(x_i, \delta_i)$. On the other hand, using differentiability of f at x_i and the above inequality, we can also deduce that

$$\left\|Df(x_i)^{-1}\left(f(x) - f(x_i)\right)\right\|_2 \geq \|x - x_i\|_2 - \varepsilon \left\|Df(x_i)^{-1}\right\|_{\mathrm{op}} \|x - x_i\|_2,$$

which, in turn, yields that

$$\|f(x) - f(x_i)\|_2 \geq \frac{1 - \varepsilon \|Df(x_i)^{-1}\|_{\mathrm{op}}}{\|Df(x_i)^{-1}\|_{\mathrm{op}}} \|x - x_i\|_2 > \varepsilon \|x - x_i\|_2$$

for all $x \in \overline{B}(x_i, \delta_i)$, where the final inequality is obtained by choosing ε sufficiently small. Therefore, for all $\delta \in (0, \delta_i]$, we have that $\left\|f|_{B(x_i,\delta)} - g|_{B(x_i,\delta)}\right\|_\infty \leq \varepsilon\delta < \operatorname{dist}(f(x_i), f(\partial B(x_i, \delta)))$. Applying Proposition 1.25, part (iii) we infer that

$$\deg(f, B(x_i, \delta), y) = \deg(g, B(x_i, \delta), y) = \operatorname{sign}(\operatorname{Jac}(g)(x_i)) = \operatorname{sign}(\operatorname{Jac}(f)(x_i))$$

for every $\delta \in (0, \delta_i]$. In the above we used formula (1.9) for the degree of mappings in $C^1(\overline{U}, \mathbb{R}^d)$. Next, we choose $\delta < \min\{\delta_1, \ldots, \delta_n\}$ sufficiently small so that the sets $(B(x_i, \delta))_{i=1}^n$ are pairwise disjoint subsets of U. In the case $n = 1$ we choose a (possibly empty) open set $U_2 \subseteq U \setminus B(x_1, \delta)$ and use $y \in \mathbb{R}^d \setminus f(\overline{U_2})$, Proposition 1.25, part (ii) and property (d2) with $U_1 = B(x_1, \delta)$ to obtain the desired result. When $n > 1$, we iteratively apply property (d2) to get

$$\begin{aligned}\deg(f, U, y) &= \sum_{i=1}^{n} \deg(f, B(x_i, \delta), y) = \sum_{i=1}^{n} \operatorname{sign}(\operatorname{Jac}(f)(x_i)) \\ &= \sum_{x \in f^{-1}(\{y\})} \operatorname{sign}(\operatorname{Jac}(f(x))).\end{aligned}$$

□

We present an additional auxiliary lemma. It says that whenever a continuous mapping in $\mathbb{R}^d$ has derivative of full rank at a point, it preserves neighbourhoods of this point. The author believes that such a statement may be a folklore; however, he did not find any reference.

Lemma 1.27. *Let $a \in \mathbb{R}^d$, $r > 0$ and $f\colon \overline{B}(a, r) \to \mathbb{R}^d$ be a continuous mapping differentiable at the point a with* $\operatorname{rank}(Df(a)) = d$. *Then there is $\delta_0 > 0$ such that for every $\delta \in (0, \delta_0]$ we have*

$$B\left(f(a), \frac{\delta}{2\left\|Df(a)^{-1}\right\|_{op}}\right) \subseteq f(B(a, \delta)).$$

Up to an affine transformation, the lemma above can be restated in the following way, as we shall see a bit later:

Lemma 1.28. *Let $\alpha \in (0, 1/3)$ and $f\colon \overline{B}(0, 1) \subset \mathbb{R}^d \to \mathbb{R}^d$ be a continuous mapping such that $\|f - \mathrm{id}\|_\infty \leq \alpha$. Then $B(f(0), (1 - 2\alpha)) \subseteq f(B(0, 1))$.*

Proof. The assumptions imply that $B(f(0), (1-2\alpha))$ is disjoint from $f(\partial B(0, 1))$. Therefore, by Proposition 1.25, part (i), the degree $\deg(f, B(0, 1), \cdot)$ is constant on the ball $B(f(0), (1 - 2\alpha))$. By Proposition 1.25, part (iii), we infer that

$$\deg(f, B(0, 1), f(0)) = \deg(\mathrm{id}, B(0, 1), f(0)) = 1,$$

since $\operatorname{dist}(f(0), f(\partial B(0, 1))) \geq 1-2\alpha > \alpha \geq \|f - \mathrm{id}\|_\infty$. The lemma follows from Proposition 1.25, part (ii), which implies that every point of $B(f(0), (1 - 2\alpha))$ has to be included in $f(B(0, 1))$. □

Now we present the proof of Lemma 1.27 using Lemma 1.28:

Proof of Lemma 1.27. Without loss of generality, we assume that $a = 0$ and that $f(0) = 0$. We write $z(y) := f(y) - Df(0)(y)$. For every $y \in \mathbb{R}^d$ we have $\|z(y)\|_2 \in o(\|y\|_2)$. We pick $\beta > 0$ small, whose precise value will be set later, and choose δ_0 small enough so that $\|z(y)\|_2 \leq \beta \|y\|_2$ for every $y \in B(0, \delta_0)$.

Since the linear mapping $Df(0)$ has full rank, its inverse $Df(0)^{-1}$ is well defined and has a finite norm. We see that $(Df(0)^{-1} \circ f)(y) - y = Df(0)^{-1}(z(y))$.

We fix $\delta \in (0, \delta_0]$ and write $g_\delta(y) := \frac{1}{\delta}(Df(0)^{-1} \circ f)(\delta y)$. The mapping g_δ is defined on the ball $\overline{B}(0,1)$ and continuous. We also get that

$$\|g_\delta - \mathrm{id}\|_\infty \le \frac{\beta\delta \left\|Df(0)^{-1}\right\|_{\mathrm{op}}}{\delta} = \beta \left\|Df(0)^{-1}\right\|_{\mathrm{op}}.$$

We set $\beta := \frac{1}{4\|Df(0)^{-1}\|_{\mathrm{op}}}$. From Lemma 1.28 (with $\alpha = 1/4$) it follows that $g_\delta(B(0,1)) \supseteq B\Big(0, \frac{1}{2}\Big)$; hence, we infer that $(Df(0)^{-1} \circ f)(B(0,\delta)) \supseteq B\Big(0, \frac{\delta}{2}\Big)$. It is not hard to observe that $\inf_{x \in S^{d-1}} \|Df(0)(x)\|_2 = \frac{1}{\|Df(0)^{-1}\|_{\mathrm{op}}}$. Consequently, $f(B(0,\delta)) \supseteq B\bigg(0, \frac{\delta}{2\|Df(0)^{-1}\|_{\mathrm{op}}}\bigg)$. $\square$

At one place we will need a special case of the multiplication theorem for the degree of a composition of two continuous mappings (see, e.g., Fonseca and Gangbo [24, Thm. 2.10]):

Theorem 1.29 (A special case of the Multiplication Theorem)**.** *Let $U \subseteq \mathbb{R}^d$ be open. Moreover, let $f\colon U \to \mathbb{R}^d$ and $g\colon f(U) \to \mathbb{R}^d$ be continuous, injective mappings. Then for every $y \in g \circ f(U)$ it holds that*

$$\deg(g \circ f, U, y) = \deg(g, f(U), y) \cdot \deg(f, U, g^{-1}(y)).$$

Bilipschitz decomposition of Lipschitz regular mappings

Our main goal in this section is to show that Lipschitz regular mappings in Euclidean spaces decompose into bilipschitz mappings in a nice way. For reader's convenience, we restate the main theorem we are going to prove here:

Theorem 1.30 (Bonk and Kleiner [5])**.** *Let $U \subseteq \mathbb{R}^d$ be open and $f\colon \overline{U} \to \mathbb{R}^d$ be Lipschitz regular. Then there are disjoint open sets $(A_n)_{n\in\mathbb{N}}$ such that $\bigcup_{n\in\mathbb{N}} A_n$ is dense in $\overline{U}$ and such that $f|_{A_n}$ is bilipschitz with the lower bilipschitz constant $b = b(\mathrm{Reg}(f))$.*

As we said before, the theorem above can be easily deduced combining Theorem 3.4 and Lemma 4.2 of Bonk and Kleiner [5]. We will present an easier and shorter (but less general) proof than that of Bonk and Kleiner [5].

Before we prove Theorem 1.30, let us put it briefly into context. For a general Lipschitz mapping $h\colon \mathbb{R}^d \to \mathbb{R}^d$ it is known that one can obtain a different bilipschitz decomposition using Sard's theorem; see, e.g., Federer [22, Lem. 3.2.2]. One can start with sets

$$\bigg\{x \in \mathbb{R}^d\colon\ Dh(x)^{-1} \text{ exists}, \left\|Dh(x)^{-1}\right\|_{\mathrm{op}} \le k \text{ and } \forall y \in B\Big(x, \frac{1}{k}\Big)$$
$$\|h(y) - h(x) - Dh(x)(y-x)\|_2 \le \frac{\|x-y\|_2}{2k}\bigg\}$$

defined for every $k \in \mathbb{N}$ and then cut these sets into pieces of diameter less than $1/k$ forming a decomposition $(A_n)_{n=1}^\infty$. Then Sard's theorem implies that $\mathcal{L}\Big(h\Big(\mathbb{R}^d \setminus \bigcup_{n\in\mathbb{N}} A_n\Big)\Big) = 0$. When compared to the decomposition established in Theorem 1.30, the difference is that the sets A_n are not necessarily open, the lower bilipschitz constant of each $h|_{A_n}$ may depend on n and $\bigcup_{n=1}^{\infty} A_n$ need not be a large subset of the domain in any sense.

If the decomposition that was just described is applied to a Lipschitz regular mapping, the resulting sets A_n occupy almost all of the domain, since the set $N(f)$ has a full measure in the domain. But the sets A_n still need not be open. The fact that for Lipschitz regular mappings it is possible to ensure the openness of bilipschitz pieces A_n will be of crucial importance to us.

The first quantitative version of the decomposition using Sard's theorem was provided by David [12, Prop. 1] for general Lipschitz mappings $f\colon \mathbb{R}^d \to \mathbb{R}^m$, where $m \geq d$. David shows that for any ball $B \subset \mathbb{R}^d$, if $\mathcal{L}(f(B))$ is large in measure, then B contains a set E large in measure such that $f|_E$ is bilipschitz. When applied to a Lipschitz regular mapping f, using the measure-preserving property expressed in Lemma 1.21, the condition that $\mathcal{L}(f(B))$ is large in measure is satisfied automatically; for this version of David's result, see David and Semmes [15, Thm. 4.1].

A well-known result of Jones [30] provides another quantitative version of the decomposition for Lipschitz mappings $I^d \to \mathbb{R}^m$. In the decomposition of Jones as well as that of David the bilipschitz pieces may have empty interior.

The presented proof of Theorem 1.30 can be divided into three parts. The first one is to find, for any given open set in the domain, an open subset on which the given mapping is almost injective (this notion is formalised below). The second part is to show that a Lipschitz regular, almost injective map on an open set is injective and the third part is to prove that a Lipschitz regular, injective map on an open set with a convex image is bilipschitz. In each of these steps the proof relies on the Lipschitz regularity of the mapping in question.

Definition 1.31. *We say that a mapping $h\colon A \subseteq \mathbb{R}^d \to \mathbb{R}^d$ is* almost injective *if there is a set $B \subseteq A$ such that $\mathcal{L}(A \setminus B) = 0$ and $h|_B$ is injective.*

As advertised above, we begin by showing that a Lipschitz regular mapping is almost injective on some open set:

Lemma 1.32. *Let $U \subseteq \mathbb{R}^d$ be non-empty and open and $f\colon U \to \mathbb{R}^d$ be Lipschitz regular. Then there is a non-empty open set $V \subseteq U$ such that $f|_V$ is almost injective and $f(V)$ is an open ball.*

Proof. The proof relies heavily on the special properties of the set $N(f)$ (see Definition 1.24 and the discussion beneath). Thanks to Observation 1.23 we know that $|f^{-1}(\{y\})|$ is finite for every $y \in f(U)$. Pick $y \in f(N(f))$ with $|f^{-1}(\{y\})|$ maximal and set $\{x_1, \dots, x_k\} = f^{-1}(\{y\}) \subseteq N(f)$. We choose pairwise disjoint, open balls $C_1, \dots, C_k$ in U centred at $x_1, \dots, x_k$, respectively. By Lemma 1.27, there is a non-empty, open ball $G \subseteq \bigcap_{i=1}^{k} f(C_i)$ centred at y. Hence, by the choice of y, the mapping f is injective on each set of the form $C_i \cap f^{-1}(G) \cap N(f)$, for $i \in [k]$. Since $N(f) \cap C_i$ occupies almost all of C_i, any $f|_{C_i \cap f^{-1}(G)}$ is almost injective. $\square$

As the next step, we use topological degree to show that whenever a Lipschitz regular mapping is almost injective on an open set U, it is injective on U.

Lemma 1.33. *Let $U \subseteq \mathbb{R}^d$ be an open set, $f\colon \overline{U} \to \mathbb{R}^d$ be a Lipschitz regular, almost injective mapping. Then $f|_U$ is injective.*

Proof. Suppose not. This means we can find two points $x_1 \neq x_2$ in U such that $y := f(x_1) = f(x_2) \in f(U)$. We pick two disjoint balls B_1, B_2 in U centred at x_1, x_2, respectively, whose boundaries do not intersect the set $f^{-1}(\{y\})$, which is finite by Observation 1.23. We may then choose $\delta > 0$ sufficiently small so that $B(y,\delta) \subseteq \mathbb{R}^d \setminus (f(\partial B_1) \cup f(\partial B_2))$.

By Proposition 1.25, part (i) the degree $\deg_i := \deg(f, B_i, \cdot)$ is constant on $B(y,\delta)$ for $i = 1,2$. If for both $i = 1,2$ we have $\deg_i|_{B(y,\delta)} \not\equiv 0$, then by Proposition 1.25, part (ii) every point in $B(y,\delta)$ has a preimage in both B_1 and B_2, which is impossible. Hence, say, $\deg_1|_{B(y,\delta)} \equiv 0$. Since $N(f)$ is dense in B_1, there are points of $f(N(f))$ in $f(B_1) \cap B(y,\delta)$. Any such point has at least two preimages in B_1 by Proposition 1.26; again, this is a contradiction. □

The third step towards the proof of Theorem 1.30 is to show that a Lipschitz regular, injective mapping with a convex image is bilipschitz.

Lemma 1.34. *Let $U \subseteq \mathbb{R}^d$ be an open set and $f\colon U \to \mathbb{R}^d$ be an injective, Lipschitz regular mapping such that $f(U)$ is convex. Then f is bilipschitz with lower bilipschitz constant at least $\frac{1}{2\operatorname{Reg}(f)^2}$.*

We note that a very similar statement with the same proof also appears in Bonk and Kleiner [5, Lemma 4.2]. For reader's convenience, we include its short proof here as well.

Proof. By Invariance of Domain (see Theorem 1.8) the mapping f is a homeomorphism onto its image.

For every two distinct points $x, y \in U$ we consider the line segment $\overline{f(x)f(y)} \subset f(U)$ connecting their images. Its preimage under f, we denote it by $\gamma(x,y) := f^{-1}\big(\overline{f(x)f(y)}\big)$, is a curve with endpoints x and y. By Lipschitz regularity (see Definition 1.12), the curve $\gamma(x,y)$ can be covered by at most $\operatorname{Reg}(f)$ balls of radius $\operatorname{Reg}(f)\,\|f(y)-f(x)\|_2$. Consequently, the distance between x and y cannot be larger than $2\operatorname{Reg}(f)^2\,\|f(y)-f(x)\|_2$. □

Finally, we have gathered all the ingredients needed for the proof of Theorem 1.30.

Proof of Theorem 1.30. We start with a countable basis $(U_n)_{n\in\mathbb{N}}$ for the subspace topology on U. By a consecutive application of Lemmas 1.32, 1.33 and 1.34 we get a collection of open sets $(V_n)_{n\in\mathbb{N}}$ such that for every $n \in \mathbb{N}$ we have $V_n \subseteq U_n$ and that $f|_{V_n}$ is bilipschitz with lower bilipschitz constant $b = 1/(2\operatorname{Reg}(f)^2)$.

Now we set $A_1 := V_1$ and inductively define $A_n := V_n \setminus \bigcup_{j=1}^{n-1} \overline{A_j}$. By construction, the set $\bigcup_{n=1}^{\infty} A_n$ is dense in U, and hence, also in $\overline{U}$. □

Remark. The first two steps described above, which comprise of Lemmas 1.32 and 1.33, may be replaced by an application of [5, Theorem 3.4]. Bonk and Kleiner [5] work with much more general mappings; instead of assuming that $f\colon \overline{U} \subseteq \mathbb{R}^d \to \mathbb{R}^d$ is Lipschitz regular, they only require that f is continuous and that there is some constant $C > 0$ such that $|f^{-1}(\{y\})| \leq C$ for all $y \in \mathbb{R}^d$. The latter condition is referred to as 'bounded multiplicity'. Moreover, the domain $\overline{U}$ may be replaced by any compact metric space X with the property that every non-empty open subset of X has topological dimension d.

The argument we presented is different to that of Bonk and Kleiner [5]. However, a key aspect of both proofs appears to be finding points x in the domain such that $f(x)$ is an interior point of the image $f(O)$ for every neighbourhood O of x. The most difficult part of Bonk and Kleiner's argument is to show that such points exist. However, for Lipschitz regular mappings we can easily find many such points using almost everywhere differentiability of Lipschitz mappings, the regularity condition and Lemma 1.27. Indeed, note that all points in the set $N(f)$ have this property. Therefore, the author believes that the argument presented here may be a more accessible approach to Theorem 1.30 for the special case where the mappings considered are Lipschitz regular.

Using Theorem 1.30 we prove Theorem 1.35, which we later use to resolve Feige's question (Question 1.4). That is, we deduce that a Lipschitz regular mapping on an open set can be expressed, on some open subset of the image, as a sum of bilipschitz homeomorphisms. Again, we restate the theorem we are going to prove here:

Theorem 1.35. *Let $U \subseteq \mathbb{R}^d$ be non-empty and open and $f\colon \overline{U} \to \mathbb{R}^d$ be a Lipschitz regular mapping. Then there exist a non-empty open set $T \subseteq f(\overline{U})$, $N \in [\mathrm{Reg}(f)]$ and pairwise disjoint open sets $W_1, \dots, W_N \subseteq \overline{U}$ such that $f^{-1}(T) = \bigcup_{i=1}^{N} W_i$ and $f|_{W_i}\colon W_i \to T$ is a bilipschitz homeomorphism for each i with the lower bilipschitz constant $b = b(\mathrm{Reg}(f))$.*

Proof. Let $(A_n)_{n=1}^{\infty}$ be the open sets from the conclusion of Theorem 1.30 applied to the mapping f. Let $y \in f(\overline{U})$ be such that the number

$$N = N_y := \left|\left\{n \in \mathbb{N}\colon f^{-1}(\{y\}) \cap A_n \neq \emptyset\right\}\right|$$

is maximal. Note that $N \in [\mathrm{Reg}(f)]$ by Observation 1.23. Choose $\beta \in \mathbb{N}^N$ such that

$$y \in f(A_n) \quad \Leftrightarrow \quad n \in \{\beta_1, \beta_2, \dots, \beta_N\}\,.$$

Set $T = \bigcap_{i=1}^{N} f(A_{\beta_i})$ and note that T is an open set containing y. We claim that $f^{-1}(T) \subseteq \bigcup_{i=1}^{N} A_{\beta_i}$. Assuming that this claim is valid we may define the desired sets $(W_i)_{i=1}^{N}$ by $W_i := f^{-1}(T) \cap A_{\beta_i}$ for each $i \in [N]$.

Thus the proof can be completed by verifying the earlier claim, that is, by proving that $f^{-1}(T) \subseteq \bigcup_{i=1}^{N} A_{\beta_i}$. Let $z \in f^{-1}(T)$. Using that the union $\bigcup_{n=1}^{\infty} A_n$ is dense in $\overline{U}$, we may find sequences $(\alpha_k)_{k=1}^{\infty} \subseteq \mathbb{N}$ and $(z_k)_{k=1}^{\infty} \subseteq \overline{U}$ with $z_k \in A_{\alpha_k}$ such that $z_k \to z$. But then $f(z_k) \to f(z) \in T$ and so we may choose $K \geq 1$ sufficiently large so that $f(z_k) \in T$ whenever $k \geq K$. By the choice of y we have that

$$f(A_n) \cap T \neq \emptyset \quad \Leftrightarrow \quad n \in \{\beta_1, \dots, \beta_N\}\,.$$

Thus we conclude that $\alpha_k \in \{\beta_1, \dots, \beta_N\}$ for all $k \geq K$ and $z = \lim_{k\to\infty} z_k \in \bigcup_{i=1}^{N} \overline{A_{\beta_i}}$.

If $z \in \partial A_{\beta_i}$ for some $i \in [N]$, then we may choose $x \in A_{\beta_i}$ such that $f(x) = f(z)$. However, this contradicts the fact that f is bilipschitz on A_{β_i}, and therefore also bilipschitz on $\overline{A_{\beta_i}}$. We conclude that $z \in \bigcup_{i=1}^{N} A_{\beta_i}$. □

Optimality of Theorem 1.30

The remainder of the current section is devoted to discussion of limits and optimality of Theorem 1.30. The content here is independent of Feige's question (Question 1.4).

A natural question that could have come to reader's mind after reading about various bilipschitz decomposition theorems for Lipschitz regular mappings may have been the following:

Question. *Can we hope for any control of the measure of the bilipschitz pieces in a bilipschitz decomposition of Lipschitz regular mappings if one requires the pieces being open? For example, can we hope for any control of the measure of the set $\bigcup_{n=1}^{\infty} A_n$ given by the conclusion of Theorem 1.30?*

The answer to the previous question is no: the decomposition from Theorem 1.30 cannot be strengthened in this way for a general Lipschitz regular mapping; below we will show that this is unavoidable. However, in a special case that a Lipschitz regular mapping f satisfies $\operatorname{Reg}(f) \leq 2$, we can provide a stronger bilipschitz decomposition; namely, the bilipschitz pieces A_n, in addition to the conclusions of Theorem 1.30, can cover almost all of the domain.

Lemma 1.36. *Let $U \subseteq \mathbb{R}^d$ be a bounded, open set with $\mathcal{L}(\partial U) = 0$ and $f\colon \overline{U} \to \mathbb{R}^d$ be a Lipschitz regular mapping with $\operatorname{Reg}(f) \leq 2$. Then there exist pairwise disjoint, open sets $(A_n)_{n=1}^{\infty}$ in $\overline{U}$ such that $\mathcal{L}(\overline{U} \setminus \bigcup_{n=1}^{\infty} A_n) = 0$ and $f|_{A_n}$ is bilipschitz with lower bilipschitz constant $b = b(\operatorname{Reg}(f))$.*

Proof. From Observation 1.23 we know that every point $y \in f(\overline{U})$ has either one or two preimages. Since $\mathcal{L}(\partial U) = 0$, the set $f(N(f)) \setminus f(\partial U)$ has full measure in $f(\overline{U})$ by Corollary 1.22 (Luzin's property (N)). Let $y \in f(N(f)) \setminus f(\partial U)$. Using Lemma 1.27, we may choose $r > 0$ sufficiently small so that $B(y,r) \subseteq f(\overline{U}) \setminus f(\partial U)$.

If $\deg(f, U, y) \equiv 1 \pmod 2$, then Proposition 1.26 implies that y has exactly one preimage. Using Proposition 1.25, part (i), we deduce that the same is true of all points $y' \in f(N(f)) \cap B(y,r)$. Thus the mapping $f\colon f^{-1}(B(y,r)) \to B(y,r)$ is almost injective. We may now apply Lemma 1.33 and then Lemma 1.34 to conclude that $f|_{f^{-1}(B(y,r))}$ is bilipschitz with lower bilipschitz constant $\frac{1}{2\operatorname{Reg}(f)^2}$.

On the other hand, if $\deg(f, U, y) \equiv 0 \pmod 2$, then y must have two distinct preimages $x_1, x_2 \in N(f)$. Let B_1, B_2 be disjoint balls with $x_1 \in B_1$ and $x_2 \in B_2$. From Lemma 1.27 we deduce that $f(B_1) \cap f(B_2)$ contains a non-empty open ball G containing the point y. Then every point in G has exactly one preimage in each of the balls B_1 and B_2. Hence $f|_{f^{-1}(G) \cap B_i}$ is injective for $i = 1, 2$ and, applying Lemma 1.34, we conclude that these mappings are also bilipschitz with lower bilipschitz constant $\frac{1}{2\operatorname{Reg}(f)^2}$.

In the above we established that for every point $y \in f(N(f)) \setminus f(\partial U)$ there is an open ball B containing y such that $f^{-1}(B)$ decomposes precisely as the union of at most two sets on which f is bilipschitz with lower bilipschitz constant $\frac{1}{2\operatorname{Reg}(f)^2}$. The collection of all such balls forms a Vitali cover of $f(N(f)) \setminus f(\partial U)$, so we can apply the Vitali covering theorem (see, e.g., Mattila [41, Thm. 2.2, p. 26]) to extract a countable, pairwise disjoint subcollection $(B_n)_{n=1}^{\infty}$ which covers almost all of the set $f(N(f)) \setminus f(\partial U)$, and so almost all of $f(U)$. The desired sets

A_n, verifying the statement of the lemma, can now be defined as the connected components of the sets $f^{-1}(B_n)$. $\square$

On the other hand, for every $\varepsilon > 0$ we provide an example of a regular mapping $f\colon I^d \to \mathbb{R}^d$ with $\operatorname{Reg}(f) = 3$ and the following property: the set of points x such that there is an open neighbourhood of x on which f is injective has measure at most ε. Consequently, for Lipschitz regular mappings f with $\operatorname{Reg}(f) \geq 3$ we cannot hope for any control of the measure of the bilipschitz pieces A_n if we insist on A_n being open (since any bilipschitz mapping is necessarily injective).

Example 1.37. *For any $\varepsilon > 0$ there is a $(3, \sqrt{d})$-regular mapping $f\colon I^d \to \sqrt{d}I^d$ and a set $X \subset I^d$ with the following properties:*

(i) $\left\|f - \sqrt{d}\,\mathrm{id}\right\|_\infty \leq \varepsilon$.

(ii) $\mathcal{L}(X) \geq 1 - \varepsilon$.

(iii) For every $x \in X$ and every $\delta > 0$ the mapping f is not injective on the ball $B(x, \delta)$. Moreover, there are disjoint, non-empty, open balls $U_1, U_2 \subseteq B(x, \delta)$ such that $\frac{1}{\sqrt{d}} f|_{U_i}$ is an isometry for $i = 1, 2$ and $f(U_1) = f(U_2)$.

Proof. We give a proof for the case $d = 1$. The example for $d \geq 1$ can easily be constructed from this: if $f\colon I \to \mathbb{R}$ is the example for the case $d = 1$ with an appropriate choice of ε, then the function $h\colon I^d \to \sqrt{d}I^d$ defined by

$$h(x_1, x_2, \dots, x_d) = (\sqrt{d}f(x_1), \sqrt{d}x_2, \dots, \sqrt{d}x_d), \qquad \text{for } (x_1, x_2, \dots, x_d) \in I^d$$

verifies Example 1.37 for general $d \geq 1$.

Given a point $a \in (0, 1)$ and $c > 0$, we will denote by $F_{a,c}$ the interval $[a, a + 3c]$. Next, we define a 1-Lipschitz function $g(a, c)\colon I \to I$ that makes two folds on $F_{a,c}$ in a sense; see Figure 1.4.

More precisely, we let

$$g(a,c)(x) := \begin{cases} x & \text{if } x \leq a + c \\ 2a + 2c - x & \text{if } x \in [a + c, a + 2c] \\ x - 2c & \text{if } x \geq a + 2c. \end{cases}$$

We will now summarise various properties of the function $g(a, c)$ which will be needed in the following construction. It is clear that $g(a, c)$ is 1-Lipschitz and $\|g(a, c) - \mathrm{id}\|_\infty \leq 2c$. Moreover, g isometrically maps each of the three subintervals $[a+(i-1)c, a+ic]$ of $F_{a,c}$, for $i \in [3]$, onto the same interval $[a, a+c]$. Denoting by J_1, J_2 the two components of the set $I \setminus F_{a,c}$ we further point out that the sets $g(a, c)(J_1)$, $g(a, c)(J_2)$ and $g(a, c)(F_{a,c})$ are pairwise disjoint subsets of I, and that g restricted to each J_i is a translation. Therefore, for any interval $U \subseteq g(a, c)(I)$, the preimage $g(a, c)^{-1}(U)$ is an isometric copy of U whenever U does not intersect $g(a, c)(F_{a,c})$, and $g(a, c)^{-1}(U)$ may be covered by 3 intervals of length $\mathcal{L}(U)$ whenever U intersects $g(a, c)(F_{a,c})$.

Let $X \subseteq I$ be a fat Cantor set[13] with $\mathcal{L}(X) \geq 1 - \varepsilon$ and $(A_n)_{n=1}^\infty$ be an enumeration of the components of $I \setminus X$. In what follows we will use the fact

[13] See, e.g., Mattila [41].

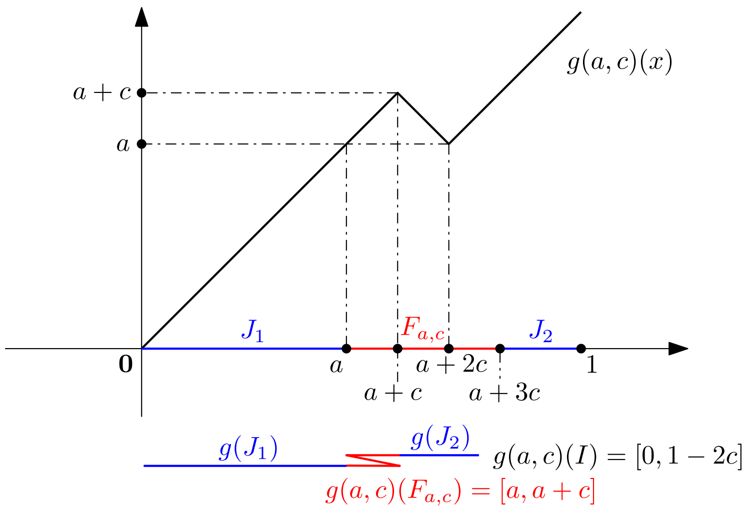

Figure 1.4: The function $g(a,c)$ and the interval $F_{a,c}$ together with their images.

that every neighbourhood of a given point $x \in X$ contains some of the intervals $(A_n)_{n=1}^{\infty}$. The idea of the construction is to 'pleat' inside each of the intervals A_n using mappings of the form $g(a,c)$ defined above; see Figure 1.5.

Now we describe the construction more formally. We start with $f_0 := \mathrm{id}$. Let a_n be a midpoint of the interval A_n. For $n \in \mathbb{N}$ we write $g_n := g(f_{n-1}(a_n), c_n)$ and $f_n := g_n \circ f_{n-1}$, where $c_n > 0$ are chosen small enough with respect to several constraints, which will be described during the course of the construction. Then we define f as the limit of f_n.

The first requirement on c_n is that $F_{a_n,c_n} \subset A_n$. Second, in order for f to be well-defined, we want to choose c_n so that the sequence $(f_n)_{n=1}^{\infty}$ is Cauchy. We have already observed that $\|g_n - \mathrm{id}\|_\infty \le 2c_n$. Thus, choosing $c_n \le \frac{\varepsilon}{2^{n+1}}$, we get that $\|f_n - f_{n-1}\|_\infty \le 2c_n \le \frac{\varepsilon}{2^n}$ and that f is well-defined. Moreover, f clearly satisfies condition (i).

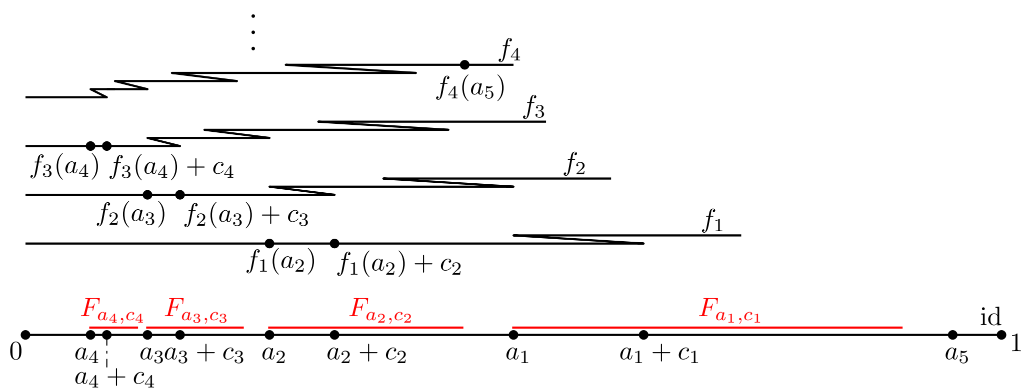

Figure 1.5: The first four steps in the construction of the functions $(f_n)_{n=1}^{\infty}$.

To see that f is $(3,1)$-regular, note that f is 1-Lipschitz because f is the uniform limit of a sequence of 1-Lipschitz functions. However, obtaining the regularity estimate is a bit more tricky.

Given any open interval $U \subseteq f(I)$, we have

$$f^{-1}(U) \subseteq f_n^{-1}(\overline{B}(U, \|f_n - f\|_\infty), \qquad \text{for } n \in \mathbb{N}.$$

Imagine, for the time being, that the latter set can be covered by 3 intervals of length $D(\mathcal{L}(U) + 2\|f_n - f\|_\infty)$ for some $D < 3$. Then letting $n \to \infty$ we deduce that $f^{-1}(U)$ can be covered by 3 *closed* intervals of length $D\mathcal{L}(U)$, which in turn can be covered by 3 *open* intervals of length $3\mathcal{L}(U)$.

Therefore, we fix a strictly increasing sequence of numbers $D_n \in [1, 3)$ such that $D := \sup D_n < 3$ and show that c_n can be chosen so that the following holds true: for every $n \in \mathbb{N}$ and every open interval $U \subset I$ the set $f_n^{-1}(U)$ can be covered by 3 intervals of length $D_n\mathcal{L}(U)$.

For $n = 0$ the condition is clearly satisfied by any $D_0 \geq 1$. For a general $n \in \mathbb{N}$ we will distinguish three cases. If U is disjoint from $f_n(F_{a_n,c_n})$, then $f_n^{-1}(U) = f_{n-1}^{-1}(g_n^{-1}(U))$ is the preimage under f_{n-1} of a translation of U, which can be covered by 3 intervals of length $D_{n-1}\mathcal{L}(U)$ by induction, which is less than $D_n\mathcal{L}(U)$.

If U intersects the interval $f_n(F_{a_n,c_n})$, but is disjoint from $\bigcup_{i=1}^{n-1} f_n(F_{a_i,c_i})$, then $f_n^{-1}(U)$ is a translation of $g_n^{-1}(U)$, which can be covered by 3 intervals of length $\mathcal{L}(U)$, as was already noted in the discussion of the properties of $g(a, c)$ above.

We are left with the option that U intersects $f_n(F_{a_n,c_n})$ and also the set $\bigcup_{i=1}^{n-1} f_n(F_{a_i,c_i})$. However, since the intervals in $(A_n)_{n=1}^\infty$ are pairwise disjoint, this means that $\mathcal{L}(U)$ must be quite large; namely, $\mathcal{L}(U) \geq \frac{\mathcal{L}(A_n)}{2} - 3c_n$, since a_n is the midpoint of A_n. On the other hand, the inequality $\|f_n - f_{n-1}\|_\infty \leq 2c_n$ implies that

$$f_n^{-1}(U) \subseteq f_{n-1}^{-1}(\overline{B}(U, 2c_n)).$$

By induction, the latter set can be covered by 3 intervals of length $D_{n-1}(\mathcal{L}(U) + 4c_n)$. The last quantity can be made smaller than $D_n\mathcal{L}(U)$ using the lower bound on $\mathcal{L}(U)$ and choosing c_n small enough. This finishes the proof that the function f is $(3, 1)$-regular.

By construction, the function f is not injective on any of the intervals $F_{a_n,c_n} \subset A_n$, but it maps each of the three subintervals $[a_n + (j-1)c_n, a_n + jc_n]$, for $j \in [3]$, isometrically onto the same interval. Since every neighbourhood of any point of X contains some of the intervals $(A_n)_{n=1}^\infty$, this verifies condition (iii). □

There is another question that can come to reader's mind. Where can we put Lipschitz regular mappings on the imaginary scale between bilipschitz and Lipschitz mappings? Are they closer to general Lipschitz mappings or rather to bilipschitz ones?

We can show that a typical 1-Lipschitz mapping, in the sense of the Baire Category Theorem (see, e.g., Munkres [43, Thm. 48.2]), is not injective on any open subset of the domain, and hence, in the light of Theorem 1.30, a typical 1-Lipschitz mapping is not Lipschitz regular. Before we state and proof the result, we present an auxiliary lemma.

Lemma 1.38. *Let $U_1, U_2 \subseteq \mathbb{R}^d$ be disjoint, non-empty open balls and $f \colon \overline{U_1} \cup \overline{U_2} \to \mathbb{R}^d$ be a continuous mapping such that $f|_{U_i}$ is a homeomorphism for $i = 1, 2$ and $f(U_1) = f(U_2)$. Then there exists $s > 0$ such that for any continuous mapping $h \colon \overline{U_1} \cup \overline{U_2} \to \mathbb{R}^d$ with $\|h - f\|_\infty < s$ we have $h(U_1) \cap h(U_2) \neq \emptyset$.*

Proof. Fix $y \in f(U_1) = f(U_2)$. Property (d1) of the degree together with Theorem 1.29 applied to $(f|_{U_i})^{-1} \circ f|_{U_i} = \mathrm{id}$ implies that $\deg(f, U_i, y) \in \{-1, 1\}$ for $i = 1, 2$.

Let $s > 0$ be small enough so that $B(y, s) \subseteq f(U_1) = f(U_2)$ and $h\colon \overline{U_1} \cup \overline{U_2} \to \mathbb{R}^d$ be a continuous mapping with $\|h - f\|_\infty < s \leq \mathrm{dist}(y, f(\partial U_i))$ for $i = 1, 2$. Then we may apply Proposition 1.25, part (iii) to obtain $\deg(h, U_i, y) = \deg(f, U_i, y) \in \{-1, 1\}$ for $i = 1, 2$. Finally, we use Proposition 1.25, part (ii) to deduce that $y \in h(U_1) \cap h(U_2) \neq \emptyset$. $\square$

Proposition 1.39. *Let $\mathcal{V}$ denote the complete metric space of 1-Lipschitz mappings $I^d \to \mathbb{R}^d$ equipped with the metric induced by the supremum norm. Then the set of all 1-Lipschitz mappings which are injective on some non-empty, open subset of I^d is a meagre subset of $\mathcal{V}$.*

Proof. Let us write $\mathcal{B}$ for a countable base of the topology on I^d consisting of open balls. Moreover, for every $D \in \mathcal{B}$ we denote by $\mathcal{I}(D)$ the subset of $\mathcal{V}$ consisting of mappings that are injective on D. It is sufficient to show that the set $\mathcal{I}(D)$ forms a nowhere dense subset of $\mathcal{V}$ for every $D \in \mathcal{B}$.

Let $D = B(u, r) \in \mathcal{B}$, $g \in \mathcal{I}(D)$ and $\eta > 0$. To verify that $\mathcal{I}(D)$ is nowhere dense we will find $g' \in \mathcal{V}$ and $s > 0$ such that $\|g' - g\|_\infty < \eta$ and $B(g', s) \cap \mathcal{I}(D)$ is empty.

Choose $\varepsilon < \min\{\mathcal{L}(B(u, r/2)), r/2, \eta\}$ and let $f\colon I^d \to \sqrt{d}I^d$ and $X \subseteq I^d$ be given by Example 1.37. Then $g \circ \frac{1}{\sqrt{d}}f \in \mathcal{V}$ and using part (i) of Example 1.37 we infer $\left\|g \circ \frac{1}{\sqrt{d}}f - g\right\|_\infty \leq \left\|\frac{1}{\sqrt{d}}f - \mathrm{id}\right\|_\infty \leq \varepsilon < \eta$. By the choice of ε, there exists $x \in B(u, r/2) \cap X$. By Example 1.37, part (iii) there are disjoint, non-empty, open balls $U_1, U_2 \subset B(u, r/2)$ such that $\frac{1}{\sqrt{d}}f|_{U_i}$ is an isometry for $i = 1, 2$ and $\frac{1}{\sqrt{d}}f(U_1) = \frac{1}{\sqrt{d}}f(U_2) =: G$. Note that $\left\|\frac{1}{\sqrt{d}}f - \mathrm{id}\right\|_\infty \leq \varepsilon < r/2$ implies that $G \subseteq B(u, r) = D$. Thus, $g|_G$ is injective and, by Invariance of Domain (see Theorem 1.8), a homeomorphism. We have now established that $g \circ \frac{1}{\sqrt{d}}f$ maps each of the two disjoint, non-empty, open balls $U_1, U_2 \subseteq I^d$ homeomorphically onto the same open set $g(G)$. It follows from Lemma 1.38 that we can choose $s > 0$ sufficiently small so that whenever $h\colon I^d \to \mathbb{R}^d$ is a continuous mapping with $\left\|h - g \circ \frac{1}{\sqrt{d}}f\right\|_\infty < s$ the sets $h(U_1)$ and $h(U_2)$ have non-empty intersection, implying that h is not injective. For s chosen as above, we have $B\left(g \circ \frac{1}{\sqrt{d}}f, s\right) \cap \mathcal{I}(D) = \emptyset$. $\square$

Yet another question that a curious reader may ask is whether Lipschitz regular mappings can be characterised as Lipschitz mappings admitting a bilipschitz decomposition as in Theorem 1.30.

However, this turns out not to be the case. It is easy to construct an example with infinitely many overlapping images of bilipschitz pieces. But even more is true: it is possible to construct an injective 1-Lipschitz function on the unit interval that has a decomposition as in Theorem 1.30, but, at the same time, is not Lipschitz regular. An example f is given by the formula

$$f(x) = \int_0^x g(t)dt, \qquad x \in I,$$

where $g\colon I \to I$ is any positive, bounded, measurable function which is constant and equal to one on a dense collection of open subintervals of I and not a.e. bounded away from zero.

1.3 Geometric properties of bilipschitz mappings

Bilipschitz mappings of a Euclidean space $\mathbb{R}^d$ transform volume according to the formula $\mathcal{L}(f(A)) = \int_A |\mathrm{Jac}(f)|\, \mathrm{d}\mathcal{L}$ (see Theorem 1.11). In this section we establish that bilipschitz mappings cannot transform volume too wildly. In some sense we show that sufficiently fine grids of cubes must witness 'continuity' of the volume transform. This in turn places rather restrictive conditions on the Jacobian of a bilipschitz mapping, which we will exploit in Section 1.4 in order to find non-realisable densities. The work presented in this section is an interpretation of the construction of Burago and Kleiner [8], which is modified in various ways, leading to some extensions of the results in [8]. Critically for the present solution of Feige's question, Burago and Kleiner's construction is adapted so that it treats multiple bilipschitz mappings simultaneously. In the light of Theorem 1.30 and Theorem 1.35 established for Lipschitz regular mappings in the previous section, this will make Burago and Kleiner's techniques applicable to Lipschitz regular mappings.

Notation. In this section we will often use numbers and vectors together. To help the reader distinguish between vectors and numbers, we typeset vectors in bold.

We write $\mathbf{e}_1, \dots, \mathbf{e}_d$ for the standard basis of $\mathbb{R}^d$ and $\mathbf{0}$ for the origin in $\mathbb{R}^d$. For $\lambda > 0$ we let $\mathcal{Q}^d_\lambda$ denote the standard tiling of $\mathbb{R}^d$ by cubes of sidelength λ and vertices in the set $\lambda\mathbb{Z}^d$. We call a family of cubes *tiled* if it is a subfamily of $\mathcal{Q}^d_\lambda$ for some $\lambda > 0$. We say that two cubes $S, S' \in \mathcal{Q}^d_\lambda$ are $\mathbf{e}_1$-*adjacent* if $S' = S + \lambda\mathbf{e}_1$. For mappings $h\colon \mathbb{R}^d \to \mathbb{R}^k$ we denote by $h^{(1)}, \dots, h^{(k)}$ the co-ordinate functions of h.

The main result of the present section is the following lemma:

Lemma 1.40. *Let $d, k \in \mathbb{N}$ with $d \geq 2$, $L \geq 1$ and $\eta, \zeta \in (0,1)$. Then there exists $r = r(d, k, L, \eta, \zeta) \in \mathbb{N}$ such that for every non-empty open set $U \subseteq \mathbb{R}^d$ there exist finite tiled families $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_r$ of cubes contained in U with the following properties:*

1. *For each $1 \leq i < r$ and each cube $S \in \mathcal{S}_i$*

$$\mathcal{L}\left(S \cap \bigcup_{j=i+1}^{r} \bigcup \mathcal{S}_j\right) \leq \eta\mathcal{L}(S).$$

2. *For any k-tuple $(h_1, \dots, h_k)$ of L-bilipschitz mappings $h_j\colon U \to \mathbb{R}^d$ there exist $i \in [r]$ and $\mathbf{e}_1$-adjacent cubes $S, S' \in \mathcal{S}_i$ such that*

$$\left| \fint_S |\mathrm{Jac}(h_j)| - \fint_{S'} |\mathrm{Jac}(h_j)| \right| \leq \zeta$$

for all $j \in [k]$.

Statement 1 expresses that each collection of cubes $\mathcal{S}_{i+1}$ is much finer than the previous collection $\mathcal{S}_i$. The inequality of statement 2 can be interpreted geometrically as stating that the volume of the image of the cube S under h_j is very close to the volume of the image of its neighbour S'. Put differently, we may rewrite the inequality of statement 2 in the following form:

$$|\mathcal{L}(h_j(S)) - \mathcal{L}(h_j(S'))| \leq \zeta \mathcal{L}(S).$$

It is possible to assemble Lemma 1.40 using predominantly arguments contained in the article of Burago and Kleiner [8]. However, Burago and Kleiner do not state any version of Lemma 1.40 explicitly and to prove Lemma 1.40 it is not sufficient to just take some continuous part of their argument. One needs to inspect their whole proof in detail and work considerably to put together all of the pieces correctly. Therefore, we present a complete proof of Lemma 1.40 in which we introduce some new elements. The proof of Lemma 1.40 requires some preparation and will be given later in this section.

Remark. Variants of the Burago–Kleiner [8] construction with additional details have been employed in a pure discrete setting in the works of Garber [26], Magazinov [37] and Cortez and Navas [10].

Auxiliary lemmas

Lying behind all of the results of the present section is a simple property of Lipschitz mappings of an interval: If $[0, c] \subseteq \mathbb{R}$ is an interval and a Lipschitz mapping $h\colon [0, c] \to \mathbb{R}^n$ stretches the endpoints 0, c almost as much as its Lipschitz constant allows, then it is intuitively clear that the mapping h is close to affine. The next dichotomy can be thought of as a 'discretised' version of this statement:

Lemma 1.41. *Let $L \geq 1$ and $\varepsilon > 0$. Then there exist parameters*

$$M = M(L, \varepsilon) \in \mathbb{N}, \qquad \varphi = \varphi(L, \varepsilon) > 0$$

such that for all $c > 0$, $n \in \mathbb{N}$, $N \in \mathbb{N}$, $N \geq 2$ and all L-Lipschitz mappings $h\colon [0, c] \to \mathbb{R}^n$ at least one of the following statements holds:

1. *There exists a set $\Omega \subset [N-1]$ with $|\Omega| \geq (1-\varepsilon)(N-1)$ such that for all $i \in \Omega$ and for all $x \in \left[\frac{(i-1)c}{N}, \frac{ic}{N}\right]$*

$$\left\| h\left(x + \frac{c}{N}\right) - h(x) - \frac{1}{N}(h(c) - h(0)) \right\|_2 \leq \frac{c\varepsilon}{N}.$$

2. *There exists $z \in \frac{c}{NM}\mathbb{Z} \cap \left[0, c - \frac{c}{NM}\right]$ such that*

$$\frac{\left\| h\left(z + \frac{c}{NM}\right) - h(z) \right\|_2}{\frac{c}{NM}} > (1 + \varphi) \frac{\|h(c) - h(0)\|_2}{c}.$$

In the lemma above, statement 1 expresses in a discrete way that h is close to affine: after partitioning the interval $[0, c]$ into N subintervals of equal length this statement asserts that h looks like an affine mapping on nearly all pairs of adjacent subintervals. Statement 2 is a discrete formulation of the condition that the Lipschitz constant of h is not almost realised by the endpoints 0, c.

Remark. *The statement of Lemma 1.41 involves plethora of parameters and the oncoming Lemma, which is its multi-dimensional version for bilipschitz mappings, is even more complicated. To help the reader make any sense of it, the* meaning *of all the parameters involved is always the same in all subsequent lemmas in this section. The author would like to explain here informally the intended meaning of all the parameters used.*

The parameters L, d, c, ε and N are chosen by the user of the lemma. The meaning of L and d is obvious: they stand for the Lipschitz/bilipschitz constant of h and the dimension of the ambient space, respectively. The role of the parametr c is to set the scale. Another approach would be to simply fix $c = 1$, say, and then argue that the results presented in this section are invariant under a proportional scaling. The number ε quantifies what it means for h to be 'close' to affine. The parameter N determines the desired 'resolution'; the user of the lemma cuts the domain into even pieces on which the mapping h is desired to be 'close to affine' in some sense. In the subsequent lemmas, the value of N will be restricted from below, but never from above.

The remaining parameters are determined by the lemma(s) and not to be chosen by the user. In the case of Lemma 1.41 we are left with the parameters M and φ. The former describes how many (regularly placed) points inside each piece *have to be checked for the failure of statement 2 in order to deduce the validity of statement 1 given the desired precision ε. The latter quantifies how far from realising Lipschitz constant on the endpoints of the interval $[0, c]$ the mapping h has to be in case of failure of statement 1.*

Note that it is crucial that M and φ do not depend on N and the mapping h; they depend only the Lipschitz/bilipschitz constant of h and ε.

We now formulate a multi-dimensional version of Lemma 1.41; see Figure 1.6. We consider thin cuboids in $\mathbb{R}^d$ of the form $[0, c] \times [0, c/N]^{d-1}$ and prove that when such a cuboid is sufficiently thin, that is, when N is sufficiently large, then the one-dimensional statement for L-Lipschitz mappings $f : [0, c] \to \mathbb{R}^n$ given in Lemma 1.41 can, in a sense, be extended to L-bilipschitz mappings $f : [0, c] \times [0, c/N]^{d-1} \to \mathbb{R}^n$.

Lemma 1.42. *Let $d \in \mathbb{N}$, $L \geq 1$ and $\varepsilon > 0$. Then there exist parameters*

$$M = M(d, L, \varepsilon) \in \mathbb{N}, \quad \varphi = \varphi(d, L, \varepsilon) \in (0, 1), \quad N_0 = N_0(d, L, \varepsilon) \in \mathbb{N}$$

such that for all $c > 0$, $N \in \mathbb{N}$, $N \geq N_0$ and all L-bilipschitz mappings

$$h : [0, c] \times [0, c/N]^{d-1} \to \mathbb{R}^n$$

at least one of the following statements holds:

1. *There exists a set $\Omega \subset [N-1]$ with $|\Omega| \geq (1-\varepsilon)(N-1)$ such that for all $i \in \Omega$ and for all $\mathbf{x} \in \left[\frac{(i-1)c}{N}, \frac{ic}{N}\right] \times \left[0, \frac{c}{N}\right]^{d-1}$*
$$\left\| h\left(\mathbf{x} + \frac{c}{N}\mathbf{e}_1\right) - h(\mathbf{x}) - \frac{1}{N}(h(c\mathbf{e}_1) - h(\mathbf{0})) \right\|_2 \leq \frac{c\varepsilon}{N}. \tag{1.10}$$

2. *There exists $\mathbf{z} \in \frac{c}{NM}\mathbb{Z}^d \cap \left(\left[0, c - \frac{c}{NM}\right] \times \left[0, \frac{c}{N} - \frac{c}{NM}\right]^{d-1}\right)$ such that*
$$\frac{\left\| h\left(\mathbf{z} + \frac{c}{NM}\mathbf{e}_1\right) - h(\mathbf{z}) \right\|_2}{\frac{c}{NM}} > (1+\varphi)\frac{\|h(c\mathbf{e}_1) - h(\mathbf{0})\|_2}{c}.$$

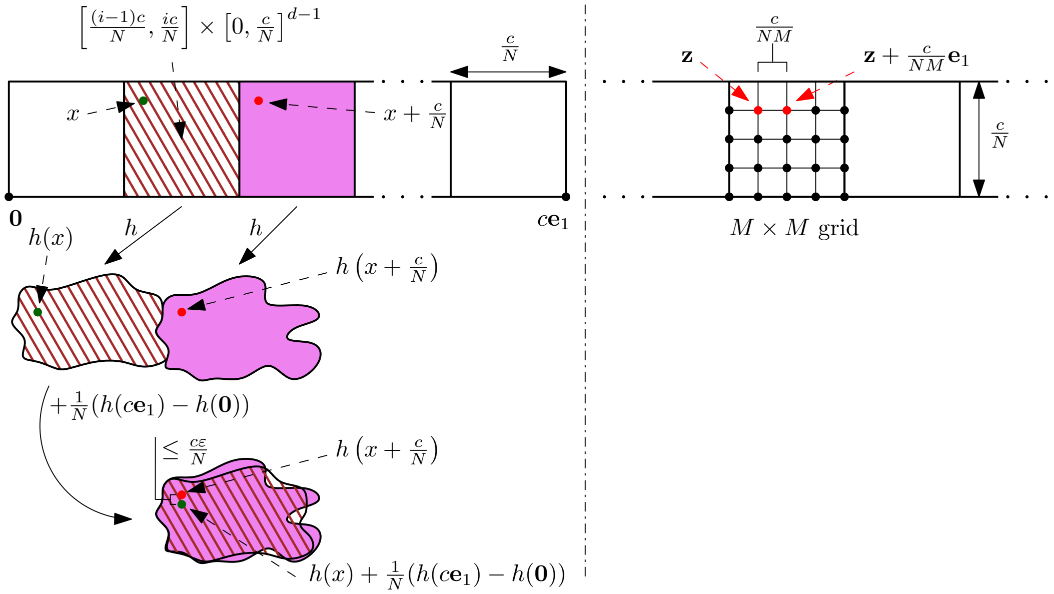

Figure 1.6: An illustration of statement 1 (left) and statement 2 (right) of Lemma 1.42. The left-hand side illustrates that h maps two neighbouring cubes to 'similar' images; after a translation by $\frac{1}{N}(h(c\mathbf{e}_1) - h(\mathbf{0}))$, the image of the left cube is pointwise at least $\frac{c\varepsilon}{N}$-close to the image of its neighbouring cube.

The proof of the one-dimensional statement Lemma 1.41 presented in this section follows Burago and Kleiner [8] closely, but a new induction argument to deduce Lemma 1.42 from Lemma 1.41 is used. This is to expose clearly that the property of bilipschitz mappings established in Lemma 1.42 is of a one-dimensional nature.

Since the proofs of Lemmas 1.41 and 1.42 are technical, we postpone their proofs until the end of this section and present first how Lemma 1.42 is going to be used.

To begin with, we show that whenever statement 1 of Lemma 1.42 holds for a bilipschitz mapping h into $\mathbb{R}^d$, there are adjacent cubes S_i and S_{i+1} whose images under h have almost the same measure. Eventually this will lead to conclusion 2 of Lemma 1.40.

Lemma 1.43. *Let $L \geq 1$, $\varepsilon \in (0, 1/2L)$, $d \in \mathbb{N}$ and $N_0 = N_0(d, L, \varepsilon)$ be given by the conclusion of Lemma 1.42. Let $N \geq N_0$, $c > 0$, $h\colon [0, c] \times [0, c/N]^{d-1} \to \mathbb{R}^d$ be an L-bilipschitz mapping, $i \in [N-1]$ and suppose that h satisfies inequality* (1.10) *on $S_i := \left[\frac{(i-1)c}{N}, \frac{ic}{N}\right] \times \left[0, \frac{c}{N}\right]^{d-1}$. Then*

$$|\mathcal{L}(h(S_i)) - \mathcal{L}(h(S_{i+1}))| \leq 2L^{d+1} d\varepsilon \mathcal{L}(S_i).$$

Proof. To simplify the notation, let us denote the cube S_i only by S in this proof. Define a translation $\phi\colon \mathbb{R}^d \to \mathbb{R}^d$ by

$$\phi(\mathbf{y}) = \mathbf{y} + \frac{1}{N}(h(c\mathbf{e}_1) - h(\mathbf{0})).$$

Let the mappings $f_1\colon S \to \mathbb{R}^d$, $f_2\colon S \to \mathbb{R}^d$ be defined by $f_1 := \phi \circ h$ and $f_2(\mathbf{x}) = h\left(\mathbf{x} + \frac{c}{N}\mathbf{e}_1\right)$. Then f_1, f_2 are both L-bilipschitz mappings of the cube

$S \in \mathcal{Q}^d_{c/N}$ which satisfy

$$\|f_1 - f_2\|_\infty \leq \frac{c\varepsilon}{N}, \tag{1.11}$$

due to (1.10). These conditions imply the bound on the difference in volume of the images $f_1(S)$ and $f_2(S)$ asserted above.

To show this, we first set up an additional notation: for a set $A \subseteq \mathbb{R}^d$ and $t > 0$ we introduce the set

$$[A]_t := \{\mathbf{x} \in A \colon \operatorname{dist}(\mathbf{x}, \partial A) \geq t\}$$

of all points in the interior of A whose distance to the boundary of A is at least t.

For all $t > 0$, using (1.11) we deduce that

$$f_1([S]_t) \subseteq \overline{B}\left(f_2([S]_t), \frac{c\varepsilon}{N}\right).$$

Moreover, by Invariance of Domain (see Theorem 1.8), f_2 preserves boundaries. In combination with the lower bilipschitz bound we see that $f_2([S]_t) \subseteq [f_2(S)]_{t/L}$. Specifically, we get that

$$\overline{B}\left(f_2([S]_t), \frac{c\varepsilon}{N}\right) \subseteq \overline{B}\left([f_2(S)]_{t/L}, \frac{c\varepsilon}{N}\right)$$

holds for all $t > 0$ as well. Combining the two inclusions above and setting $t := \frac{Lc\varepsilon}{N}$, it follows that

$$f_1\left([S]_{\frac{Lc\varepsilon}{N}}\right) \subseteq f_2(S).$$

Therefore,

$$\mathcal{L}(f_1(S)) - \mathcal{L}(f_2(S)) \leq \mathcal{L}\left(f_1\left(S \setminus [S]_{\frac{Lc\varepsilon}{N}}\right)\right) \leq L^d \mathcal{L}\left(S \setminus [S]_{\frac{Lc\varepsilon}{N}}\right),$$

where the last inequality follows from the upper Lipschitz bound on f_1 (see equation (1.3) in 'Background and notation').

The Lebesgue measure of the set $S \setminus [S]_{\frac{Lc\varepsilon}{N}}$ can be easily computed using the fact that $[S]_{\frac{Lc\varepsilon}{N}}$ is a cube of sidelength $\frac{c}{N}(1 - 2L\varepsilon)$ inside the cube S:

$$\begin{aligned}\mathcal{L}\left(S \setminus [S]_{\frac{Lc\varepsilon}{N}}\right) = \mathcal{L}(S) - \mathcal{L}\left([S]_{\frac{Lc\varepsilon}{N}}\right) &= \left(\frac{c}{N}\right)^d \left(1 - (1 - 2L\varepsilon)^d\right) \\ &\leq \left(\frac{c}{N}\right)^d 2dL\varepsilon = 2dL\varepsilon\mathcal{L}(S).\end{aligned}$$

For the inequality we use $2L\varepsilon \in (0, 1)$ and apply Bernoulli's inequality. We conclude that

$$\mathcal{L}(f_1(S)) - \mathcal{L}(f_2(S)) \leq 2dL^{d+1}\varepsilon\mathcal{L}(S).$$

Since the above argument is completely symmetric with respect to f_1 and f_2, we also have

$$\mathcal{L}(f_2(S)) - \mathcal{L}(f_1(S)) \leq 2dL^{d+1}\varepsilon\mathcal{L}(S).$$

□

Given a bilipschitz mapping $g \colon [0, c] \times [0, c/N]^{d-1} \to \mathbb{R}^n$, we now seek to repetitively apply Lemma 1.42 on smaller and smaller scales in order to, in some sense, eliminate statement 2 of the dichotomy of Lemma 1.42. Consequently, we find cubes on which g satisfies inequality (1.10) of statement 1 of Lemma 1.42 (scaled and translated copies of the sets $\left[\frac{(i-1)c}{N}, \frac{ic}{N}\right] \times \left[0, \frac{c}{N}\right]^{d-1}$). This will allow us to apply Lemma 1.43.

Strategy to eliminate statement 2 from Lemma 1.42

Let all parameters d, L, ε, M, φ, N_0, c, n and N be given by the statement of Lemma 1.42. We consider an L-bilipschitz mapping $g\colon [0,c]\times[0,c/N]^{d-1}\to\mathbb{R}^n$. If statement 2 holds for g, there is a pair of points $\mathbf{a}_1 := \mathbf{z}, \mathbf{b}_1 := \mathbf{z} + \frac{c}{NM}\mathbf{e}_1$ which the mapping g stretches by a factor $(1+\varphi)$ more than it stretches the pair $\mathbf{a}_0 := \mathbf{0}$ and $\mathbf{b}_0 := c\mathbf{e}_1$. We may now consider the restriction of g to a rescaled copy of the original cuboid $[0,c]\times[0,c/N]^{d-1}$ with vertices $\mathbf{a}_1$ and $\mathbf{b}_1$ corresponding to $\mathbf{0}$ and $c\mathbf{e}_1$, respectively. If, again, it is the case that statement 2 is valid for this mapping, then we find points $\mathbf{a}_2, \mathbf{b}_2$ inside the new cuboid which g stretches by a factor $(1+\varphi)$ more than it stretches the pair $\mathbf{a}_1$ and $\mathbf{b}_1$, and so a factor $(1+\varphi)^2$-times more than it stretches $\mathbf{a}_0$ and $\mathbf{b}_0$. The process is illustrated in Figure 1.7. We iterate this procedure as long as possible to obtain sequences $(\mathbf{a}_i)$ and $(\mathbf{b}_i)$ satisfying

$$\frac{\|g(\mathbf{b}_i)-g(\mathbf{a}_i)\|}{\|\mathbf{b}_i-\mathbf{a}_i\|} \geq (1+\varphi)^i \frac{\|g(\mathbf{b}_0)-g(\mathbf{a}_0)\|}{\|\mathbf{b}_0-\mathbf{a}_0\|} \geq \frac{(1+\varphi)^i}{L},$$

where the final bound is given by the lower bilipschitz inequality for g. It is clear now that the procedure described above cannot continue forever: otherwise, for i sufficiently large, the inequality above contradicts the L-Lipschitz condition on g. Thus, Lemma 1.42 tells us that after at most r iterations of the procedure, where $r\in\mathbb{N}$ is a number determined by d, L and ε, we must have that statement 1 is valid for an appropriate restriction of the mapping g.

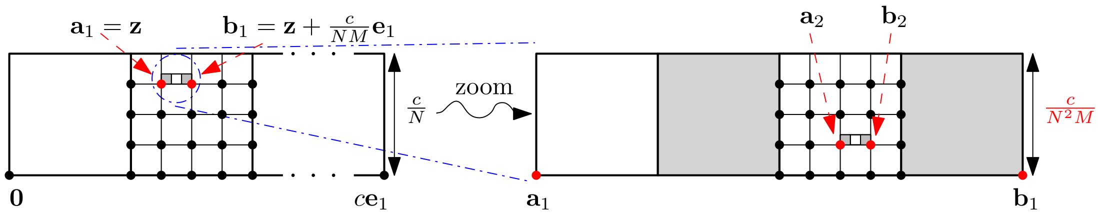

Figure 1.7: An illustration of the strategy to eliminate statement 2 from Lemma 1.42. The stretch factor of g on the points $\mathbf{a}_1$ and $\mathbf{b}_1$ is at least $(1+\varphi)$-times larger then the stretch factor of g on the points $\mathbf{0}$ and $c\mathbf{e}_1$. If statement 2 applies in the next iteration, we find two points stretched by g with factor at least $(1+\varphi)^2$-times the stretch factor on $\mathbf{0}$ and $c\mathbf{e}_1$.

Let us now present the strategy described in the sketch above formally.

Lemma 1.44. *Let $d\in\mathbb{N}$, $L\geq 1$, $\varepsilon>0$, the parameters $M = M(d,L,\varepsilon)$, and $N_0 = N_0(d,L,\varepsilon)$ be given by Lemma 1.42, $c>0$, $N\geq N_0$. Moreover, let*

$$g\colon [0,c]\times[0,c/N]^{d-1}\to\mathbb{R}^n$$

be an L-bilipschitz mapping and $c_i := \frac{c}{(NM)^{i-1}}$ for $i\in\mathbb{N}$. Then there exist a parameter $r = r(d,L,\varepsilon)\in\mathbb{N}$ and

$$\mathbf{z}_1 = \mathbf{0}, \quad \mathbf{z}_{i+1} \in c_{i+1}\mathbb{Z}^d \cap [0, c_i - c_{i+1}] \times \left[0, \frac{c_i}{N} - c_{i+1}\right]^{d-1} \qquad \textit{for } i\in[r-1]$$

such that statement 1 of Lemma 1.42 is valid for at least one *of the mappings*

$$g_i\colon [0,c_i]\times[0,c_i/N]^{d-1}\to\mathbb{R}^n,\quad i\in[r],$$

defined by

$$g_i(\mathbf{x}) := g\left(\mathbf{x}+\sum_{j=1}^{i}\mathbf{z}_j\right). \tag{1.12}$$

Proof. The appropriate condition on the parameter $r = r(d,L,\varepsilon)\in\mathbb{N}$ will be determined later in the proof. Let $\varphi=\varphi(d,L,\varepsilon)$ be the parameter given by the conclusion of Lemma 1.42. We follow the procedure described below:

Procedure 1.44.1. *Set $i=1$, $\mathbf{z}_1=\mathbf{0}$ and $g_1=g$.*

1. *If statement* 1 *of Lemma 1.42 holds for $h=g_i$ and $c=c_i$, then stop. If not, proceed to step 2.*
2. *Choose $\mathbf{z}_{i+1}\in c_{i+1}\mathbb{Z}^d\cap[0,c_i-c_{i+1}]\times\left[0,\frac{c_i}{N}-c_{i+1}\right]^{d-1}$ such that*

$$\frac{\|g_i(\mathbf{z}_{i+1}+c_{i+1}\mathbf{e}_1)-g_i(\mathbf{z}_{i+1})\|_2}{c_{i+1}}>(1+\varphi)\frac{\|g_i(c_i\mathbf{e}_1)-g_i(\mathbf{0})\|_2}{c_i} \tag{1.13}$$

   and define $g_{i+1}\colon[0,c_{i+1}]\times[0,c_{i+1}/N]^{d-1}\to\mathbb{R}^n$ by

$$g_{i+1}(\mathbf{x}) := g_i(\mathbf{x}+\mathbf{z}_{i+1}) = g\left(\mathbf{x}+\sum_{j=1}^{i+1}\mathbf{z}_j\right)$$

3. *Set $i=i+1$ and return to step 1.*

At each iteration $i\geq 1$ of Procedure 1.44.1, the conditions of Lemma 1.42 are satisfied for d, L, ε, M, φ, N_0, $c=c_i$, N and $h=g_i\colon[0,c_i]\times[0,c_i/N]^{d-1}\to\mathbb{R}^n$. Therefore, whenever the algorithm does not terminate in step 1, we have that such a point $\mathbf{z}_{i+1}$ required by step 2 exists by Lemma 1.42.

To complete the proof, it suffices to verify that Procedure 1.44.1 terminates after at most r iterations. This is clear after rewriting (1.13) in the form

$$\frac{\|g_{i+1}(c_{i+1}\mathbf{e}_1)-g_{i+1}(\mathbf{0})\|_2}{c_{i+1}}>(1+\varphi)\frac{\|g_i(c_i\mathbf{e}_1)-g_i(\mathbf{0})\|_2}{c_i}>(1+\varphi)^i\frac{1}{L},$$

where the latter inequality follows by induction and the lower bilipschitz inequality for $g=g_1$. Since each g_i is L-Lipschitz, Procedure 1.44.1 cannot complete more than $r:=\left\lfloor\log_{(1+\varphi)}L^2\right\rfloor$ iterations. □

Proof of Lemma 1.40

We are finally ready to give a proof of the main result of the present section.

Proof of Lemma 1.40. Let $\varepsilon=\varepsilon(\zeta,d,L,k)\in(0,\zeta)$ be a parameter to be determined later in the proof, $M=M(d,L\sqrt{k},\varepsilon)$, $\varphi=\varphi(d,L\sqrt{k},\varepsilon)$ and $N_0=N_0(d,L\sqrt{k},\varepsilon)$ be given by the statement of Lemma 1.42. Moreover, let $N\geq N_0$ and $r=r(d,L\sqrt{k},\varepsilon)\in\mathbb{N}$ be given by the conclusion of Lemma 1.44. We impose additional conditions on ε and N in the course of the proof.

Let $U \subseteq \mathbb{R}^d$ be a non-empty open set. Since the conclusion of Lemma 1.40 is invariant under translation of the set $U \subseteq \mathbb{R}^d$, we may assume that $\mathbf{0} \in U$ and choose $c > 0$ such that

$$[0, c] \times [0, c/N]^{d-1} \subseteq U.$$

We can now define the families of cubes $\mathcal{S}_1, \ldots, \mathcal{S}_r$, making use of the sequence $c_i = \frac{c}{(NM)^{i-1}}$ defined in Lemma 1.44; see also Figure 1.8.

Definition 1.45. *For each $i \in [r]$ we define the family $\mathcal{S}_i \subseteq \mathcal{Q}_{c_i/N}$ as the collection of all cubes of the form*

$$\sum_{j=1}^{i} \mathbf{z}_j + \left[\frac{(l-1)c_i}{N}, \frac{lc_i}{N}\right] \times \left[0, \frac{c_i}{N}\right]^{d-1}$$

where $\mathbf{z}_1 = 0$, $\mathbf{z}_{j+1} \in c_{j+1}\mathbb{Z}^d \cap [0, c_j - c_{j+1}] \times \left[0, \frac{c_j}{N} - c_{j+1}\right]^{d-1}$ for each $j \geq 1$ and $l \in [N]$.

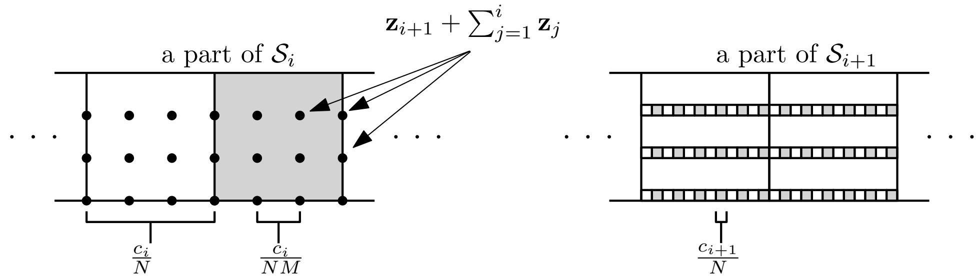

Figure 1.8: Left: two cubes from the family $\mathcal{S}_i$ with points of the form $\sum_{j=1}^{i+1} \mathbf{z}_j$, where $\mathbf{z}_j$ are fixed for $j = 1, \ldots, i$. Right: the resulting part of the family $\mathcal{S}_{i+1}$.

Let us now verify that the above defined families $\mathcal{S}_1, \ldots, \mathcal{S}_r$ satisfy condition 1 in the statement of Lemma 1.40. It is immediate from Definition 1.45 that

$$\bigcup \mathcal{S}_r \subseteq \bigcup \mathcal{S}_{r-1} \subseteq \ldots \subseteq \bigcup \mathcal{S}_1.$$

Thus, given $1 \leq i < r$ and $S \in \mathcal{S}_i$, we have that

$$S \cap \bigcup_{j=i+1}^{r} \bigcup \mathcal{S}_j \subseteq S \cap \bigcup \mathcal{S}_{i+1}.$$

Note that any cube in the collection $\mathcal{S}_{i+1}$ has the form

$$\mathbf{w} + \left[\frac{(l-1)c_{i+1}}{N}, \frac{lc_{i+1}}{N}\right] \times \left[0, \frac{c_{i+1}}{N}\right]^{d-1}$$

for some $\mathbf{w} \in c_{i+1}\mathbb{Z}^d$ and $l \in [N]$. Since $S \in \mathcal{Q}_{c_i/N}$ and $c_i/N = Mc_{i+1}$, such a cube can only intersect S in a set of positive Lebesgue measure when $\mathbf{w} \in S$. Therefore, the number of cubes in $\mathcal{S}_{i+1}$ that can intersect $S \in \mathcal{Q}_{c_i/N}$ in a set of positive Lebesgue measure is bounded above by

$$N \left| c_{i+1}\mathbb{Z}^d \cap S \right| \leq N \left(\frac{c_i/N}{c_{i+1}} + 1 \right)^d = N(M+1)^d.$$

It follows that

$$\mathcal{L}\left(S \cap \bigcup_{j=i+1}^{r} \bigcup \mathcal{S}_j\right) \le N(M+1)^d \left(\frac{c_{i+1}}{N}\right)^d = \frac{(M+1)^d}{M^d N^{d-1}} \left(\frac{c_i}{N}\right)^d \le \eta \mathcal{L}(S).$$

where, in the above, we use $c_{i+1} = c_i/NM$ and $\mathcal{L}(S) = (c_i/N)^d$ and prescribe that N is sufficiently large so that the inequality holds. Thus, statement 1 is satisfied.

Turning now to statement 2, we consider a k-tuple $(h_1, \dots, h_k)$ of L-bilipschitz mappings $h_i \colon U \to \mathbb{R}^d$ and define a mapping $g \colon U \to \mathbb{R}^{kd}$ co-ordinate-wise by

$$g^{((i-1)d+j)}(\mathbf{x}) := h_i^{(j)}(\mathbf{x})$$

for $i \in [k]$ and $j \in [d]$. It is straightforward to verify that g is $L\sqrt{k}$-bilipschitz. The conditions of Lemma 1.44 are now satisfied for d, $L\sqrt{k}$ in place of L, ε, M, N_0, c, $n = kd$ and $g \colon [0,c] \times [0, c/N]^{d-1} \to \mathbb{R}^{kd}$.

Let $p \in [r]$ be such that statement 1 of Lemma 1.42 holds for the mapping $g_p \colon [0, c_p] \times [0, c_p/N]^{d-1} \to \mathbb{R}^{kd}$, which was defined by equation (1.12) (existence of such mapping is ensured by the conclusion of Lemma 1.44). Moreover, Lemma 1.44 provides us with points $\mathbf{z}_1, \dots, \mathbf{z}_p \in \mathbb{R}^d$.

Let Ω be given by the assertion of Lemma 1.42, statement 1 for g_p. Recall that the co-ordinate functions of the mapping $g_p \colon [0, c_p] \times [0, c_p/N]^{d-1} \to \mathbb{R}^{kd}$ are defined by

$$g_p^{((t-1)d+s)}(\mathbf{x}) = g^{((t-1)d+s)}\left(\mathbf{x} + \sum_{j=1}^{p} \mathbf{z}_j\right) = h_t^{(s)}\left(\mathbf{x} + \sum_{j=1}^{p} \mathbf{z}_j\right)$$

for $t \in [k]$, $s \in [d]$. Therefore, each $h_{t,p} \colon [0, c_p] \times [0, c_p/N]^{d-1} \to \mathbb{R}^d$ defined by $h_{t,p}(\mathbf{x}) := h_t\left(\mathbf{x} + \sum_{j=1}^{p} \mathbf{z}_j\right)$ for $t \in [k]$ satisfies inequality 1.10 (with $h = h_{t,p}$) on S_i, for each $i \in \Omega$.

We fix $i \in \Omega$ and impose the condition $\varepsilon < \frac{1}{2L}$ on ε. Then, for each $t \in [k]$, the conditions of Lemma 1.43 are satisfied with $L\sqrt{k}$ in place of L, ε, d, N, $c = c_p$, $h = h_{t,p}$ and i. Hence,

$$|\mathcal{L}(h_{t,p}(S_i)) - \mathcal{L}(h_{t,p}(S_{i+1}))| \le 2(L\sqrt{k})^{d+1} d\varepsilon \mathcal{L}(S_i) \le \zeta \mathcal{L}(S_i),$$

when we prescribe that $\varepsilon \le \frac{\zeta}{2(L\sqrt{k})^{d+1} d}$. Set $S = \sum_{j=1}^{p} \mathbf{z}_j + S_i$ and $S' = \sum_{j=1}^{p} \mathbf{z}_j + S_{i+1}$. It is clear upon reference to Definition 1.45 that S and S' are $\mathbf{e}_1$-adjacent cubes belonging to the family $\mathcal{S}_p$. Moreover, we have $h_t(S) = h_{t,p}(S_i)$ and $h_t(S') = h_{t,p}(S_{i+1})$ for all $t \in [k]$. Therefore S and S' verify statement 2 of Lemma 1.40 for the k-tuple $(h_1, \dots, h_k)$. This completes the proof of Lemma 1.40. $\square$

Proofs of auxiliary Lemmas 1.41 and 1.42

To complete the section, it remains to prove Lemmas 1.41 and 1.42. Since Lemma 1.41 serves as the base case for induction in the prove of Lemma 1.42, we present the proof of Lemma 1.41 first:

Proof of Lemma 1.41. Let $M \in \mathbb{N}$ and $\varphi \in (0,1)$ be parameters to be determined later in the proof. Let $c > 0$, $n \in \mathbb{N}$, $N \in \mathbb{N}$ and $h\colon [0,c] \to \mathbb{R}^n$ be an L-Lipschitz mapping. The assertion of the Lemma holds for h if and only if the assertion holds for $\rho \circ h$, where $\rho\colon \mathbb{R}^n \to \mathbb{R}^n$ is any distance preserving transformation. Therefore, we may assume that $h(0) = (0,0,\dots,0)$ and $h(c) = (A,0,\dots,0)$ where $A := \|h(c) - h(0)\|_2$. Since h is L-Lipschitz, we see that $A \leq Lc$.

Assume that statement (2) does not hold for h. In other words, we have that

$$\frac{\left\|h\left(x + \frac{c}{NM}\right) - h(x)\right\|_2}{\frac{c}{NM}} \leq (1+\varphi)\frac{A}{c} \tag{1.14}$$

for all $x \in \frac{c}{NM}\mathbb{Z} \cap \left[0, c - \frac{c}{NM}\right]$. We complete the proof by verifying that statement (1) holds for h.

Since h is only Lipschitz and not necessarily bilipschitz, it may happen that $A = 0$. We need to treat this case separately. Thus, at first, we assume that $h(0) = h(c)$. Using (1.14), we get that $h(z) = h(0)$ for every $z \in \frac{c}{NM}\mathbb{Z} \cap [0,c]$. For any $x \in \left[0, c - \frac{c}{N}\right]$ we can find $z \in \frac{c}{NM}\mathbb{Z} \cap [0,c]$, $z \leq x$, such that $|x - z| \leq \frac{c}{NM}$. This, however, implies that

$$\begin{aligned}&\left\|h\left(x + \frac{c}{N}\right) - h(x) - \frac{1}{N}(h(c) - h(0))\right\|_2 \\ &= \left\|h\left(x + \frac{c}{N}\right) - h\left(z + \frac{c}{N}\right) + h\left(z + \frac{c}{N}\right) - h(z) + h(z) - h(x)\right\|_2 \leq \frac{2Lc}{NM}.\end{aligned}$$

The last quantity is at most $\frac{c\varepsilon}{N}$ provided $M \geq \frac{2L}{\varepsilon}$, which verifies statement (1) with $\Omega := [N-1]$.

From now on, we assume that $A > 0$. For later use, we point out that (1.14) implies

$$\|h(b) - h(a)\|_2 \leq (1+\varphi)\frac{A}{c}\|b - a\|_2 \tag{1.15}$$

whenever $a, b \in \frac{c}{NM}\mathbb{Z} \cap [0,c]$. Let $S_i = \left[\frac{(i-1)c}{N}, \frac{ic}{N}\right]$, for $i \in [N]$. We introduce an additional parameter $t = t(L,\varepsilon) \in (\varphi, 1)$ whose value will be determined later in the proof. Moreover, let us define

$$P := \left\{x \in \frac{c}{NM}\mathbb{Z} \cap \left[0, c - \frac{c}{N}\right] : h^{(1)}\left(x + \frac{c}{N}\right) - h^{(1)}(x) > \frac{(1-t)A}{N}\right\}.$$

For $x \in P$ we have

$$\left|h^{(1)}\left(x + \frac{c}{N}\right) - h^{(1)}(x) - \frac{A}{N}\right| \leq \frac{tA}{N}.$$

This inequality follows from the definition of P, the inequality (1.15) and $t > \varphi$. For the remaining co-ordinate functions we have

$$\sum_{i=2}^{n}\left|h^{(i)}\left(x + \frac{c}{N}\right) - h^{(i)}(x)\right|^2 \leq \frac{(1+\varphi)^2A^2}{N^2} - \frac{(1-t)^2A^2}{N^2} \leq \frac{4tA^2}{N^2}.$$

Combining the two inequalities above and using $t < 1$ we deduce that

$$\left\|h\left(x + \frac{c}{N}\right) - h(x) - \frac{1}{N}(h(c) - h(0))\right\|_2 \leq \frac{\sqrt{t^2 + 4t}A}{N} \leq \frac{\sqrt{5t}A}{N} \leq \frac{\sqrt{5t}Lc}{N} \tag{1.16}$$

for every $x \in P$.

Let $\Gamma \subset [0,1]$ be a maximal c/N-separated subset of $\frac{c}{NM}\mathbb{Z} \cap \left[0, c - \frac{c}{N}\right] \setminus P$ and let $x_1, \ldots, x_{|\Gamma|}$ be the elements of Γ. Then the intervals $\left(\left[x_i, x_i + \frac{c}{N}\right]\right)_{i=1}^{|\Gamma|}$ can only intersect in the endpoints. Therefore the set $[0,c] \setminus \bigcup_{i=1}^{|\Gamma|} \left[x_i, x_i + \dfrac{c}{N}\right]$ is a finite union of intervals with endpoints in $\frac{c}{NM}\mathbb{Z} \cap [0,c]$ and with total length $c - \frac{|\Gamma| c}{N}$. Using $\Gamma \cap P = \emptyset$ and (1.15) we deduce that

$$A = h^{(1)}(c) - h^{(1)}(0) \leq |\Gamma| \frac{(1-t)A}{N} + (1+\varphi)\frac{A}{c}\left(c - \frac{|\Gamma|\, c}{N}\right).$$

Since $A > 0$, we can rearrange this inequality to obtain

$$|\Gamma| \leq \frac{\varphi}{\varphi + t} N \leq \frac{2\varphi}{\varphi + t}(N-1),$$

where, for the last inequality, we apply $N \geq 2$. Since the intervals S_i can share only the endpoints, it follows that the set $\frac{c}{NM}\mathbb{Z} \cap \left[0, c - \frac{c}{N}\right] \setminus P$ can intersect at most $\frac{4\varphi}{\varphi+t}(N-1)$ intervals S_i. Letting

$$\Omega := \left\{ i \in [N-1] \colon \frac{c}{NM}\mathbb{Z} \cap S_i \subseteq P \right\}$$

we deduce that $|\Omega| \geq \left(1 - \frac{4\varphi}{\varphi+t}\right)(N-1)$. Moreover, for any $i \in \Omega$ and $x \in S_i$ we can find $x' \in P$ with $|x' - x| \leq \frac{c}{NM}$. This allows us to apply (1.16) to get

$$\begin{aligned}
&\left\| h\left(x + \frac{c}{N}\right) - h(x) - \frac{1}{N}(h(c) - h(0)) \right\|_2 \\
&\qquad \leq \left\| h\left(x' + \frac{c}{N}\right) - h(x') - \frac{1}{N}(h(c) - h(0)) \right\|_2 + \frac{2Lc}{NM} \leq \frac{Lc\left(\sqrt{5t} + \frac{2}{M}\right)}{N}.
\end{aligned}$$

We are now ready to specify the parameters t, M and φ so that the inequalities obtained above verify statement 1. First, we prescribe that $t \in (0,1)$ is sufficiently small and $M \in \mathbb{N}$ is sufficiently large so that $L\left(\sqrt{5t} + \frac{2}{M}\right) < \varepsilon$. Finally we demand that $\varphi \in (0,t)$ is small enough so that $\frac{4\varphi}{\varphi+t} < \varepsilon$. $\square$

Proof of Lemma 1.42. In this proof we will sometimes add the superscript d or $d-1$ to objects such as the Lebesgue measure $\mathcal{L}$ or vectors $\mathbf{e}_i$, $\mathbf{0}$ in order to emphasize the dimension of the Euclidean space to which they correspond. For $d \geq 2$, we will express points in $\mathbb{R}^d$ in the form $\mathbf{x} = (x_1, x_2, \ldots, x_d)$. Given $\mathbf{x} = (x_1, \ldots, x_d) \in \mathbb{R}^d$ and $s \in \mathbb{R}$, we let

$$\mathbf{x} \wedge s = (x_1, \ldots, x_d, s)$$

denote the point in $\mathbb{R}^{d+1}$ formed by concatenation of $\mathbf{x}$ and s.

The case $d = 1$ is dealt with by Lemma 1.41. Let $d \geq 2$ and suppose that the statement of the lemma holds when d is replaced with $d-1$. Given $L \geq 1$ and $\varepsilon > 0$, we let $M := M(d, L, \varepsilon) \in \mathbb{N}$, $\varphi := \varphi(d, L, \varepsilon) \in (0,1)$ and $N_0 := N_0(d, L, \varepsilon) \in \mathbb{N}$ be parameters on which we impose various conditions in the course of the proof.

For now, we just prescribe that $N_0 \geq N_0(d-1, L, \theta)$, $0 < \varphi < \frac{1}{2}\varphi(d-1, L, \theta)$ and $M \in M_{d-1}\mathbb{N}$ for $M_{d-1} := M(d-1, L, \theta)$, where $\theta := \theta(d, L, \varepsilon)$ is an additional parameter to be determined later in the proof. Note that it is important to choose M as a multiple of M_{d-1} so that $\frac{c}{NM_{d-1}}\mathbb{Z} \subseteq \frac{c}{NM}\mathbb{Z}$ whenever $N \in \mathbb{N}$.

Let $c > 0$, $N \geq N_0$ and $h\colon [0, c] \times [0, c/N]^{d-1} \to \mathbb{R}^n$ be an L-bilipschitz mapping. Note that for topological reasons we must have $n \geq d$; otherwise there would be no bilipschitz mapping $\mathbb{R}^d \to \mathbb{R}^n$.

For each $s \in [0, c/N]$ we apply the induction hypothesis to the mapping $h{\wedge}s\colon [0, c] \times [0, c/N]^{d-2} \to \mathbb{R}^n$ defined by

$$h{\wedge}s(\mathbf{x}) := h(\mathbf{x} \wedge s) = h(x_1, x_2, \ldots, x_{d-1}, s).$$

Thus, we get that for each $s \in [0, c/N]$ at least one of the following statements holds:

(1_s) There exists a set $\Omega_s \subset [N-1]$ with $|\Omega_s| \geq (1-\theta)(N-1)$ such that for all $i \in \Omega_s$

$$\left\| h{\wedge}s\left(\mathbf{x} + \frac{c}{N}\mathbf{e}_1^{d-1}\right) - h{\wedge}s(\mathbf{x}) - \frac{1}{N}\left(h{\wedge}s\left(c\mathbf{e}_1^{d-1}\right) - h{\wedge}s\left(\mathbf{0}^{d-1}\right)\right)\right\|_2 \leq \frac{c\theta}{N}$$

for all $\mathbf{x} \in \left[\frac{(i-1)c}{N}, \frac{ic}{N}\right] \times \left[0, \frac{c}{N}\right]^{d-2}$.

(2_s) There exists $\mathbf{z}_s \in \frac{c}{NM_{d-1}}\mathbb{Z}^{d-1} \cap \left(\left[0, c - \frac{c}{NM_{d-1}}\right] \times \left[0, \frac{c}{N} - \frac{c}{NM_{d-1}}\right]^{d-2}\right)$ such that

$$\frac{\left\| h{\wedge}s\left(\mathbf{z}_s + \frac{c}{NM_{d-1}}\mathbf{e}_1^{d-1}\right) - h{\wedge}s(\mathbf{z}_s)\right\|_2}{\frac{c}{NM_{d-1}}} > (1+2\varphi)\frac{\left\| h{\wedge}s\left(c\mathbf{e}_1^{d-1}\right) - h{\wedge}s\left(\mathbf{0}^{d-1}\right)\right\|_2}{c}.$$

Suppose first that statement (2_s) holds for some $s \in [0, c/N]$. Then we choose a number $s' \in \frac{c}{NM}\mathbb{Z} \cap \left[0, \frac{c}{N} - \frac{c}{NM}\right]$ with $s' \leq s$ and $|s' - s| \leq \frac{c}{NM}$. Setting $\mathbf{w} = \mathbf{z}_s \wedge s'$ we note that $\mathbf{w}$ is an element of $\frac{c}{NM}\mathbb{Z}^d \cap \left[0, c - \frac{c}{NM_{d-1}}\right] \times \left[0, \frac{c}{N} - \frac{c}{NM}\right]^{d-1}$ satisfying $\|\mathbf{w} - \mathbf{z}_s \wedge s\|_2 \leq \frac{c}{NM}$. Moreover, we see that

$$\left\| h{\wedge}s\left(c\mathbf{e}_1^{d-1}\right) - h{\wedge}s\left(\mathbf{0}^{d-1}\right)\right\|_2 \geq \left\| h(c\mathbf{e}_1^d) - h(\mathbf{0}^d)\right\|_2 - \frac{2Lc}{N}.$$

We use these inequalities and the inequality of (2_s) to derive

$$
\begin{aligned}
&\left\| h\left(\mathbf{w}+\frac{c}{NM_{d-1}}\mathbf{e}_1^d\right)-h(\mathbf{w})\right\|_2\\
&\geq \left\| h\wedge s\left(\mathbf{z}_s+\frac{c}{NM_{d-1}}\mathbf{e}_1^{d-1}\right)-h\wedge s(\mathbf{z}_s)\right\|_2-\frac{2Lc}{NM}\\
&> (1+2\varphi)\left(\frac{\left\|h\wedge s\left(c\mathbf{e}_1^{d-1}\right)-h\wedge s(\mathbf{0}^{d-1})\right\|_2}{NM_{d-1}}\right)-\frac{2Lc}{NM}\\
&\geq (1+2\varphi)\left(\frac{\left\|h\left(c\mathbf{e}_1^{d}\right)-h\left(\mathbf{0}^{d}\right)\right\|_2}{NM_{d-1}}-\frac{2Lc}{N^2M_{d-1}}\right)-\frac{2Lc}{NM}\\
&\geq \left(1+2\varphi-\frac{2(1+2\varphi)Lc}{N\left\|h(c\mathbf{e}_1^d)-h(\mathbf{0}^d)\right\|_2}-\frac{2LcM_{d-1}}{M\left\|h(c\mathbf{e}_1^d)-h(\mathbf{0}^d)\right\|_2}\right)\frac{\left\|h(c\mathbf{e}_1^d)-h(\mathbf{0}^d)\right\|_2}{NM_{d-1}}\\
&\geq \left(1+2\varphi-\frac{2(1+2\varphi)L^2}{N}-\frac{2L^2M_{d-1}}{M}\right)\frac{\left\|h\left(c\mathbf{e}_1^d\right)-h\left(\mathbf{0}^d\right)\right\|_2}{NM_{d-1}}\\
&> (1+\varphi)\frac{\left\|h\left(c\mathbf{e}_1^d\right)-h\left(\mathbf{0}^d\right)\right\|_2}{NM_{d-1}}.
\end{aligned}
$$

To deduce the penultimate inequality in the sequence above we use the lower bilipschitz bound on h. In fact, this is the only place in the proof of Lemma 1.42 where we use that the mapping h is bilipschitz and not just Lipschitz. The final inequality is ensured by taking N_0 and M sufficiently large (after fixing φ). From the final lower bound obtained for the quantity $\left\|h(\mathbf{w}+\frac{c}{NM_{d-1}}\mathbf{e}_1)-h(\mathbf{w})\right\|_2$ it follows that there exists $i\in\left[\frac{M}{M_{d-1}}\right]$ such that the point $\mathbf{z}:=\mathbf{w}+\frac{(i-1)c}{NM}\mathbf{e}_1$ verifies statement 2 for h.

We may now assume that the first statement (1_s) holds for all $s\in[0,c/N]$. We complete the proof by verifying statement 1 for h. Whenever $s\in[0,c/N]$ and $\mathbf{x}\in[0,c]\times[0,c/N]^{d-2}$ satisfy the inequality of (1_s) we have that

$$
\left\|h\left((\mathbf{x}\wedge s)+\frac{c}{N}\mathbf{e}_1^d\right)-h(\mathbf{x}\wedge s)-\frac{1}{N}\left(h\left(c\mathbf{e}_1^d\right)-h\left(\mathbf{0}^d\right)\right)\right\|_2\leq\frac{c\theta}{N}+\frac{2Lc}{N^2}. \tag{1.17}
$$

From this point onwards, let R denote the cuboid $\left[0,c-\frac{c}{N}\right]\times\left[0,\frac{c}{N}\right]^{d-1}$ and

$$
A=\left\{\mathbf{x}\in R\colon \mathbf{x}\text{ satisfies (1.10) with }\varepsilon=\theta+\frac{2L}{N}\right\}.
$$

Using (1.17) and the fact that statement (1_s) holds for every $s\in[0,c/N]$ we deduce

$$
\mathcal{L}^{d-1}(A\cap\{\mathbf{x}\colon x_d=s\})\geq(1-\theta)\mathcal{L}^{d-1}(R\cap\{\mathbf{x}\colon x_d=s\})\qquad\text{for all } s\in[0,c/N].
$$

Therefore, by Fubini's theorem (see, e.g., [41, Thm. 1.14]),

$$
\mathcal{L}^d(A)\geq(1-\theta)\mathcal{L}^d(R).
$$

For each $i \in [N-1]$ we let $S_i := \left[\frac{(i-1)c}{N}, \frac{ic}{N}\right] \times \left[0, \frac{c}{N}\right]^{d-1}$. Define

$$\Omega = \left\{ i \in [N-1] \colon \mathcal{L}^d(A \cap S_i) \geq \left(1 - \sqrt{\theta}\right) \mathcal{L}^d(S_i) \right\}$$

and observe that

$$\mathcal{L}^d(A) \leq |\Omega| \frac{\mathcal{L}^d(R)}{N-1} + (N - 1 - |\Omega|) \left(1 - \sqrt{\theta}\right) \frac{\mathcal{L}^d(R)}{N-1}.$$

Combining the two inequalities derived above for $\mathcal{L}^d(A)$, we deduce

$$\frac{|\Omega|}{N-1} \geq \left(1 - \sqrt{\theta}\right).$$

Moreover, for any $i \in \Omega$ and any cube $Q \subseteq S_i$ with side length[14] $\sqrt[d]{2\sqrt{\theta}\mathcal{L}^d(S_i)}$ we have $A \cap Q \neq \emptyset$. Therefore, for any $i \in \Omega$ and any $\mathbf{x} \in S_i$ we can find $\mathbf{x}' \in A \cap S_i$ with

$$\|\mathbf{x}' - \mathbf{x}\|_2 \leq \sqrt{d}\sqrt[d]{2\sqrt{\theta}\mathcal{L}^d(S_i)} \leq 2\sqrt{d}\,\sqrt[2d]{\theta}\frac{c}{N}.$$

Using this approximation, we obtain

$$\begin{aligned}
&\left\| h\left(\mathbf{x} + \frac{c}{N}\mathbf{e}_1\right) - h(\mathbf{x}) - \frac{1}{N}(h(c\mathbf{e}_1) - h(\mathbf{0})) \right\|_2 \\
&\leq \left\| h\left(\mathbf{x} + \frac{c}{N}\mathbf{e}_1\right) - h\left(\mathbf{x}' + \frac{c}{N}\mathbf{e}_1\right) \right\|_2 \\
&+ \left\| h\left(\mathbf{x}' + \frac{c}{N}\mathbf{e}_1\right) - h(\mathbf{x}') - \frac{1}{N}(h(c\mathbf{e}_1) - h(\mathbf{0})) \right\|_2 + \|h(\mathbf{x}') - h(\mathbf{x})\|_2 \\
&\leq 2L\|\mathbf{x}' - \mathbf{x}\|_2 + \frac{c(\theta + \frac{2L}{N})}{N} \leq \frac{c(4L\sqrt{d}\,\sqrt[2d]{\theta} + \theta + \frac{2L}{N})}{N}.
\end{aligned}$$

Thus, statement 1 is verified when we prescribe that $\theta > 0$ is sufficiently small and N_0 is sufficiently large so that

$$\left(1 - \sqrt{\theta}\right) \geq 1 - \varepsilon \quad \text{and} \quad 4L\sqrt{d}\,\sqrt[2d]{\theta} + \theta + \frac{2L}{N_0} < \varepsilon.$$

$\square$

[14] We remind that the superscript d in expressions like $\mathcal{L}^d(A)$ does not mean a d-th power of the measure of A, but merely emphasises that we are using the d-dimensional Lebesgue measure.

1.4 Realisability in spaces of functions

The main objective of the present section is to prove that in some sense almost all continuous functions $\rho \in C(I^d)$ are not *Lipschitz regular realisable*, that is, do not admit a Lipschitz regular mapping $f\colon I^d \to \mathbb{R}^d$ such that

$$f_\sharp \rho\mathcal{L} = \mathcal{L}|_{f(I^d)}. \tag{1.18}$$

More precisely, we prove Theorem 1.46, which we restate for reader's convenience:

Theorem 1.46. *The set*

$$\mathcal{E} := \Big\{ \rho \in C(I^d) : \rho \text{ admits a Lipschitz regular solution to equation } (1.18) \Big\}$$

forms a σ-porous[15] *subset of the space of continuous functions $C(I^d)$ with the supremum norm.*

To recall the meaning of σ-porosity, see Definition 1.14.

Remark. To be able to work with functions $\rho \in C(I^d)$ attaining negative values as well, we extend the definition of the pushforward measure to such functions:

$$f_\sharp \rho\mathcal{L} := f_\sharp \rho^+\mathcal{L} - f_\sharp \rho^-\mathcal{L},$$

where by ρ^+, ρ^- we mean the positive and the negative part of ρ. Technically speaking, the pushforward measure is no longer a measure, but a difference of two measures[16]*. However, we will use it only in the form of* (1.18)*, that is, when the result is again a measure.*

This is only a technical tool that helps us treat functions attaining negative values properly, but it does not bring in any additional difficulty. An alternative option would be to say that, by definition, no function with negative values satisfies (1.18)*, but the statement of Theorem 1.46 would be then seemingly weaker.*

In addition to Theorem 1.46, we also provide a related results for the space L^∞. Namely, in Theorem 1.50 later in this section it will be proven that bilipschitz realisable functions form a σ-porous subset of the space $L^\infty(I^d)$. That bilipschitz non-realisable functions contain a dense G_δ subset[17] in both the set of positive continuous functions and the set of positive, L^∞-bounded, measurable functions on the unit square $[0,1]^2$, was recently proved by Viera [47], but [47] is completely independent from the present work.

Remark. We point out that there are positive, bilipschitz non-realisable densities in $C(I^d)$ which fail to be Lipschitz regular non-realisable, i.e., positive functions $\rho \in C(I^d)$ for which equation (1.18) *admits Lipschitz regular but not bilipschitz solutions $f\colon I^d \to \mathbb{R}^d$. An example may be constructed as follows.*

We split the unit cube I^d in half, distinguishing two pieces $D_1 := [0, \frac{1}{2}] \times I^{d-1}$ and $D_2 := [\frac{1}{2}, 1] \times I^{d-1}$ and write $f\colon I^d \to D_1$ for the mapping which 'folds D_2 onto D_1'. More precisely, the mapping f is defined as the identity mapping on

[15] It means that a typical continuous function is not in $\mathcal{E}$.

[16] Sometimes this is called a *signed measure* or a *charge* in the literature.

[17] That is, a countable intersection of open sets.

D_1 and as the reflection in the hyperplane $\left\{\frac{1}{2}\right\} \times \mathbb{R}^{d-1}$ on D_2. Let $\psi \in C(D_1)$ be a positive, bilipschitz non-realisable density with values in $(0,1)$. We impose an additional mild condition that ψ is constant with value $\frac{1}{2}$ inside the hyperplane $\frac{1}{2} \times \mathbb{R}^{d-1}$. The existence of such a density ψ follows easily from the d-dimensional analog of [8, Theorem 1.2] by Burago and Kleiner.

Set $\rho := \psi$ on D_1. The bilipschitz non-realisability of ρ is now already assured, no matter how we define ρ on D_2. To make ρ Lipschitz regular realisable, we define ρ on D_2 by

$$\rho(x) := 1 - \psi(f(x)).$$

The function $\rho\colon I^d \to \mathbb{R}$ is continuous and positive, whilst the mapping $f\colon I^d \to \mathbb{R}^d$ is Lipschitz regular and satisfies $f(I^d) = D_1$. Moreover, for any measurable set $S \subseteq D_1$ we have

$$\begin{aligned} f_\sharp \rho\mathcal{L}(S) &= \int_{f^{-1}(S)\cap D_1} \rho \, d\mathcal{L} + \int_{f^{-1}(S)\cap D_2} \rho \, d\mathcal{L} \\ &= \int_S \psi \, d\mathcal{L} + \int_{f^{-1}(S)\cap D_2} (1 - \psi(f(x))) \, d\mathcal{L} \\ &= \int_S \psi \, d\mathcal{L} + \int_S (1-\psi) \, d\mathcal{L} = \mathcal{L}(S), \end{aligned}$$

where, for the penultimate equation, we use the change of variables formula (see Theorem 1.11) in conjunction with the fact that f restricted to the set D_2 is an affine isometry. This verifies the Lipschitz regular realisability of ρ.

To prove Theorem 1.46, we describe the partition of the set $\mathcal{E}$ from Theorem 1.46 into a countable family of porous sets $(\mathcal{E}_{C,L,n})$. We will need the lower bilipschitz constant $b(\cdot)$ given by the conclusion of Theorem 1.35. Recall that Theorem 1.35 provides an open set T in the image of a given Lipschitz regular mapping f, on which we can express f exactly as a sum of N bilipschitz homeomorphisms. This allows for a countable decomposition of the class of Lipschitz regular mappings: we split it into classes consisting of (C,L)-Lipschitz regular mappings for which an open set O_n of a countable basis for topology is contained inside the set T whose existence is assured by Theorem 1.35. The idea behind the decomposition $(\mathcal{E}_{C,L,n})$ of $\mathcal{E}$ presented below is to translate the countable decomposition of Lipschitz regular mappings described above to the set $\mathcal{E}$ via the change of variables formula.

Now we describe everything formally. From now on, we fix $(O_n)_{n=1}^\infty$ a countable basis for topology of I^d. For $C, L, n \in \mathbb{N}$ we let $\mathcal{E}_{C,L,n}$ denote the set of all functions $\rho \in C(I^d)$ which admit $N \in [C]$, pairwise disjoint, non-empty, open sets $Y_1, \ldots, Y_N \subseteq I^d$ with $Y_1 := O_n$, an open set $V \subseteq \mathbb{R}^d$ and $(b(C), L)$-bilipschitz homeomorphisms $f_i\colon Y_i \to V$ such that

$$\rho(y) = |\mathrm{Jac}(f_1)(y)| - \sum_{i=2}^{N} \rho\left(f_i^{-1} \circ f_1(y)\right) \left|\mathrm{Jac}(f_i^{-1} \circ f_1)(y)\right| \qquad \text{for a.e. } y \in O_n. \tag{1.19}$$

Note that the basis set O_n 'generates' the diagram of bilipschitz homeomorphisms $f_i\colon Y_i \to V$ in the sense that we have $Y_i = f_i^{-1} \circ f_1(O_n)$ for $i \in [N]$ and $V = f_1(O_n)$; see Figure 1.9. However, the critical role of O_n in the definition above is to prescribe the portion of the domain I^d on which all functions $\rho \in \mathcal{E}_{C,L,n}$ have the special form given by (1.19).

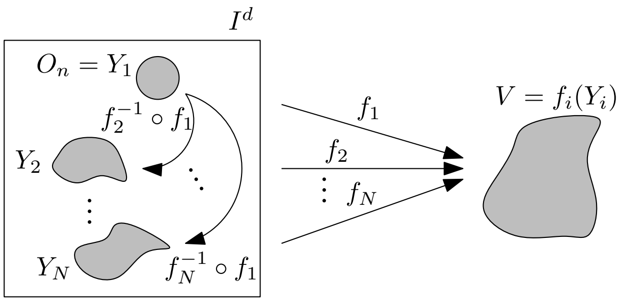

Figure 1.9: The diagram of a bilipschtz decomposition for a density $\rho \in \mathcal{E}_{C,L,n}$.

To explain the origins of equation (1.19), consider a Lipschitz regular mapping $f\colon I^d \to \mathbb{R}^d$ and the non-empty open set $T \subseteq f(I^d)$ given by the conclusion of Theorem 1.35. Because the preimage $f^{-1}(T)$ decomposes precisely as a union of N sets on which f defines a bilipschitz homeomorphism to T, a pushforward (signed) measure of the form $f_\sharp \rho \mathcal{L}$ with $\rho \in C(I^d)$ can be expressed on T as a sum of integrals involving ρ and Jacobians of N bilipschitz homeomorphisms $f_1, \dots, f_N$. Thus, whenever $f_\sharp \rho \mathcal{L} = \mathcal{L}|_{f(I^d)}$ we obtain some equation relating ρ to finitely many bilipschitz homeomorphisms and their Jacobians. We will see in the proof of lemma below that this equation has precisely the form of (1.19).

For $C, L \in \mathbb{N}$, let $\mathcal{E}_{C,L}$ denote the subset of $C(I^d)$ consisting of all functions ρ for which there exists a (C, L)-regular mapping $f\colon I^d \to \mathbb{R}^d$ solving equation (1.18). Clearly, we have $\mathcal{E} = \bigcup_{C,L \in \mathbb{N}} \mathcal{E}_{C,L}$. In the next lemma we prove that $\mathcal{E}_{C,L}$ is covered by the sets $(\mathcal{E}_{C,L,n})_{n \in \mathbb{N}}$.

Lemma 1.47. *Let $C, L \in \mathbb{N}$. Then $\mathcal{E}_{C,L} \subseteq \bigcup_{n \in \mathbb{N}} \mathcal{E}_{C,L,n}$.*

Proof. Let $\rho \in \mathcal{E}_{C,L}$ and choose a (C, L)-regular mapping $f\colon I^d \to \mathbb{R}^d$ such that $f_\sharp \rho \mathcal{L} = \mathcal{L}|_{f(I^d)}$. Let the integer $N \in [C]$ and the open sets $T \subseteq f(I^d)$ and $W_1, \dots, W_N \subseteq I^d$ be given by the conclusion of Theorem 1.35. Recall that $(O_n)_{n=1}^\infty$ is a countable basis for topology on I^d used to define $\mathcal{E}_{C,L,n}$. We choose $n \in \mathbb{N}$ such that $O_n \subseteq W_1$ and define $Y_i = f^{-1}(f(O_n)) \cap W_i$, $V = f(O_n)$ and $f_i := f|_{Y_i}\colon Y_i \to V$ for each $i \in [N]$.

To see that these choices witness that $\rho \in \mathcal{E}_{C,L,n}$, it only remains to verify equation (1.19). Note that $f^{-1}(V) = \bigcup_{i=1}^N Y_i$. Therefore, using change of variables for bilipschitz mappings (see Theorem 1.11), for every measurable set $S \subseteq V$ we infer that

$$\begin{aligned}\mathcal{L}(S) = f_\sharp \rho \mathcal{L}(S) &= \int_{f^{-1}(S)} \rho \,\mathrm{d}\mathcal{L} = \sum_{i=1}^N \int_{f^{-1}(S) \cap Y_i} \rho \,\mathrm{d}\mathcal{L} = \sum_{i=1}^N \int_{f_i^{-1}(S)} \rho \,\mathrm{d}\mathcal{L} \\ &= \sum_{i=1}^N \int_S \left(\rho \circ f_i^{-1}\right) \left|\mathrm{Jac}\left(f_i^{-1}\right)\right| \mathrm{d}\mathcal{L} = \int_S \sum_{i=1}^N \frac{\rho}{|\mathrm{Jac}(f_i)|} \circ f_i^{-1} \,\mathrm{d}\mathcal{L}.\end{aligned}$$

Since the constant function with the value one is the density of the Lebesgue measure, we conclude that

$$\sum_{i=1}^N \frac{\rho}{|\mathrm{Jac}(f_i)|} \circ f_i^{-1}(x) = 1 \qquad \text{for a.e. } x \in V.$$

Recall that the sets Y_i and V are all bilipschitz homeomorphic via the mappings $f_i\colon Y_i \to V$. Therefore, we may make the substitution $x = f_1(y)$ in the above equation after which a simple rearrangement yields

$$\rho(y) = |\mathrm{Jac}(f_1)(y)| - \sum_{i=2}^{N} \left(\frac{\rho}{|\mathrm{Jac}(f_i)|} \circ f_i^{-1} \circ f_1(y) \right) |\mathrm{Jac}(f_1)(y)| \quad \text{for a.e. } y \in O_n.$$

An application of the 'chain rule identity' for Jacobians yields (1.19). $\square$

If, for the time being, we treat the terms $\rho(f_i^{-1} \circ f_1(y))$ in (1.19) as constants, then, on the open set O_n, functions $\rho \in \mathcal{E}_{C,L,n}$ are linear combinations of at most C Jacobians of $L/b(C)$-bilipschitz mappings. The purpose of the next lemma is to provide, for given constants k and L, a function $\psi \in C(I^d)$ which is small in supremum norm, but far away from being a linear combination of k L-bilipschitz Jacobians.

Lemma 1.48. *Let $\varepsilon, \zeta \in (0,1)$, $L \geq 1$, $k \in \mathbb{N}$ and $U \subseteq I^d$ be an open set. Then there exists a function $\psi \in C(I^d)$ such that $\|\psi\|_\infty \leq \varepsilon$, $\mathrm{supp}(\psi) \subseteq U$ and for every k-tuple $(h_1, h_2, \ldots, h_k)$ of L-bilipschitz mappings $h_i\colon U \to \mathbb{R}^d$ there exist $\mathbf{e}_1$-adjacent cubes $S, S' \subseteq U$ such that*

$$\left| \fint_S |\mathrm{Jac}(h_i)| - \fint_{S'} |\mathrm{Jac}(h_i)| \right| \leq \zeta \tag{1.20}$$

for all $i \in [k]$ and

$$\left| \fint_S \psi - \fint_{S'} \psi \right| \geq \varepsilon. \tag{1.21}$$

Informally, to prove Lemma 1.48 it suffices to consider the families of tiled cubes $\mathcal{S}_1, \ldots, \mathcal{S}_r$ given by the conclusion of Lemma 1.40 applied to d, k, U, L, ζ and some very small $\eta \in (0,1)$, and to define ψ as a 'chessboard function' whose average value on $\mathbf{e}_1$-adjacent cubes makes jumps of size at least ε. From conclusion 1 of Lemma 1.40 we may essentially regard the cubes from two different families $\mathcal{S}_i$, $\mathcal{S}_j$ as pairwise disjoint; choosing η sufficiently small ensures that the values of ψ on $\bigcup \mathcal{S}_j$ have negligible impact on the average values of ψ on cubes in $\mathcal{S}_i$ for $i < j$. Below we present a formal argument.

Proof of Lemma 1.48. Let $r \in \mathbb{N}$ and the finite, tiled families $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_r$ of cubes in U be given by the conclusion of Lemma 1.40 applied to d, k, L, ζ and $\eta \in (0,1)$, where η is a parameter to be determined later in the proof. We will now define a sequence $\psi_0, \psi_1, \psi_2, \ldots, \psi_r$ of continuous functions on U. The sought after function ψ will then be defined on U by $\psi|_U := \psi_r$.

We begin by setting $\psi_0 = 0$. If $i \geq 0$ and ψ_i is already constructed, we define ψ_{i+1} as any continuous function on U with the following properties:

(i) $\psi_i = \psi_{i-1}$ outside of $\bigcup \mathcal{S}_i$.

(ii) $-\varepsilon \leq \psi_i \leq \varepsilon$.

(iii) For every cube $S \in \mathcal{S}_i$, $\fint_S \psi_i \in \{-8\varepsilon/9, 8\varepsilon/9\}$.

(iv) For every pair of $\mathbf{e}_1$-adjacent cubes $S, S' \in \mathcal{S}_i$, $\fint_S \psi_i \neq \fint_{S'} \psi_i$.

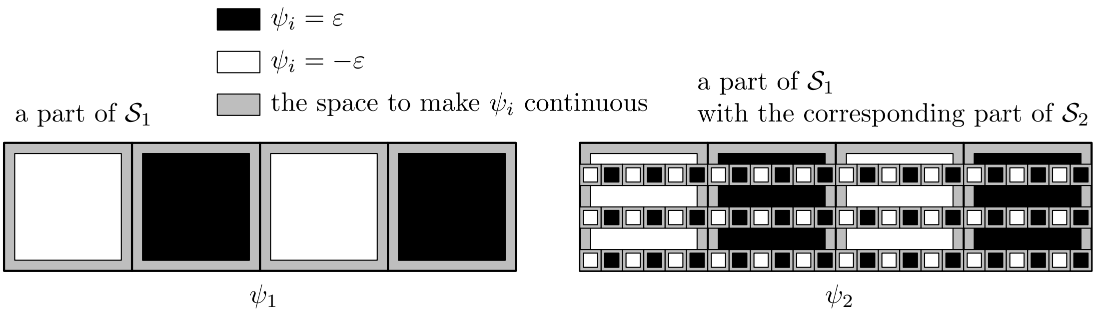

Figure 1.10: An example of the construction of the sequence of functions $\psi_1, \dots, \psi_r$.

It is clear that such a continuous function exists; see Figure 1.10 for an example. Note that conditions (iii) and (iv) just prescribe that the average values of ψ_i on the cubes in $\mathcal{S}_i$ follow a 'chessboard' pattern.

The final function ψ_r clearly satisfies $-\varepsilon \leq \psi_r \leq \varepsilon$ and $\psi_r = 0$ outside of $\bigcup_{i=1}^r \bigcup \mathcal{S}_i$. Moreover, for any $i \in [r]$ and $\mathbf{e}_1$-adjacent cubes $S, S' \in \mathcal{S}_i$ we will prove that

$$\left| \fint_S \psi_r - \fint_{S'} \psi_r \right| \geq \varepsilon. \tag{1.22}$$

Fix $i \in [r]$, $\mathbf{e}_1$-adjacent cubes $S, S' \in \mathcal{S}_i$ and combine (iii) and (iv) to obtain $\left| \fint_S \psi_i - \fint_{S'} \psi_i \right| = \dfrac{16\varepsilon}{9}$. Letting $T := S \cap \bigcup_{j=i+1}^r \bigcup \mathcal{S}_j$ and $T' := S' \cap \bigcup_{j=i+1}^r \bigcup \mathcal{S}_j$, we have from conclusion 1 of Lemma 1.40 that $\max\{\mathcal{L}(T), \mathcal{L}(T')\} \leq \eta \mathcal{L}(S)$. Moreover, condition (i) in the construction above guarantees that $\psi_r = \psi_i$ on $(S \setminus T) \cup (S' \setminus T')$. We conclude that

$$\begin{aligned} \left| \fint_S \psi_r - \fint_{S'} \psi_r \right| &\geq \left| \fint_S \psi_i - \fint_{S'} \psi_i \right| - \left| \frac{1}{\mathcal{L}(S)} \int_T (\psi_i - \psi_r) - \frac{1}{\mathcal{L}(S)} \int_{T'} (\psi_i - \psi_r) \right| \\ &\geq \frac{16\varepsilon}{9} - \frac{2 \left\| \psi_i - \psi_r \right\|_\infty \max\{\mathcal{L}(T), \mathcal{L}(T')\}}{\mathcal{L}(S)} \geq \frac{16\varepsilon}{9} - 4\varepsilon\eta. \end{aligned}$$

Thus, setting $\eta = 1/9$, we verify (1.22). To complete the proof, it now only remains to extend the function ψ to the whole of I^d by setting $\psi := 0$ outside of U so that $\operatorname{supp}(\psi) \subseteq U$. $\square$

Now we proceed to the key proof of the present section, namely the verification of porosity of the sets $(\mathcal{E}_{C,L,n})$. We actually prove that the sets $(\mathcal{E}_{C,L,n})$ possess a stronger property:

Lemma 1.49. *For every $C, L, n \in \mathbb{N}$, $\mathcal{E}_{C,L,n}$ is a porous subset of $C(I^d)$. In fact, the set $\mathcal{E}_{C,L,n}$ is porous at every point $\phi \in C(I^d)$.*

Before we begin the proof, we will outline the strategy. For given $C, L, n \in \mathbb{N}$, $\phi \in C(I^d)$ and $\varepsilon \in (0,1)$, our task is to find a function $\widetilde{\phi} \in C(I^d)$ so that $\left\| \widetilde{\phi} - \phi \right\|_\infty \leq \varepsilon$ and $B\left(\widetilde{\phi}, \alpha\varepsilon\right) \cap \mathcal{E}_{C,L,n} = \emptyset$ for some $\alpha = \alpha(C, L, n, \phi)$.

We will exploit the uniform continuity of ϕ: by prescribing at the start a sufficiently small open set $U \subseteq O_n$ we may treat ϕ as constant (relative to ε) on U and indeed on any $L/b(C)$-bilipschitz image of U (recall that $b(C)$ stands for the lower bilipschitz constant from the conclusion of Theorem 1.35). Thus, when using the condition (1.19) for functions $\rho \in \mathcal{E}_{C,L,n}$ we will always be able to treat the terms $\rho\left(f_i^{-1} \circ f_1(y)\right)$ as constant. In other words, on U we will have that all functions in $\mathcal{E}_{C,L,n}$ are linear combinations of at most C $L/b(C)$-bilipschitz Jacobians. We set $\widetilde{\phi} := \phi + \psi$ where ψ is given by the conclusion of Lemma 1.48 for $\zeta = \alpha\varepsilon$, $L' = L/b(C)$ and an appropriate choice of k. If $B\left(\widetilde{\phi}, \alpha\varepsilon\right) \cap \mathcal{E}_{C,L,n}$ is non-empty, then, up until addition by the 'constant' ϕ, the function ψ is approximately a linear combination of at most C $L/b(C)$-bilipschitz Jacobians on U. This will be incompatible with the conclusion of Lemma 1.48.

Proof of Lemma 1.49. Let $C, L, n \in \mathbb{N}$, $\phi \in C(I^d)$ and $\varepsilon \in (0,1)$. We will construct $\widetilde{\phi} \in C(I^d)$ with $\left\|\widetilde{\phi} - \phi\right\|_\infty \leq \varepsilon$ and $B\left(\widetilde{\phi}, \zeta\right) \cap \mathcal{E}_{C,L,n} = \emptyset$ for a parameter $\zeta \in (0, \varepsilon)$ to be determined later in the proof.

Using that ϕ is uniformly continuous, we may choose $\delta > 0$ sufficiently small so that

$$|\phi(y) - \phi(x)| \leq \zeta \qquad \text{whenever } y, x \in I^d \text{ and } \|y - x\|_2 \leq \delta. \tag{1.23}$$

Next, we choose an open subset $U \subseteq O_n$ with $\operatorname{diam}(U) \leq \delta b(C)/L$.

Let $\psi \in C(I^d)$ be given by the conclusion of Lemma 1.48 applied to ε, ζ, $L' = L/b(C)$, $k = C$ and U. We define the function $\widetilde{\phi} \in C(I^d)$ by

$$\widetilde{\phi} := \phi + \psi.$$

From the conclusion of Lemma 1.48 we have that $\left\|\widetilde{\phi} - \phi\right\|_\infty \leq \varepsilon$ and $\widetilde{\phi} = \phi$ outside of the set $U \subseteq O_n$. Let us now verify that $B\left(\widetilde{\phi}, \zeta\right) \cap \mathcal{E}_{C,L,n} = \emptyset$.

Let $\rho \in B\left(\widetilde{\phi}, \zeta\right)$ and suppose for a contradiction that $\rho \in \mathcal{E}_{C,L,n}$. Choose $N \in [C]$, pairwise-disjoint, non-empty, open sets $Y_1, \ldots, Y_N \subseteq I^d$, $V \subseteq \mathbb{R}^d$ and $(b(C), L)$-bilipschitz homeomorphisms $f_i \colon Y_i \to V$ witnessing that $\rho \in \mathcal{E}_{C,L,n}$. By the choice of ψ and Lemma 1.48 there exist $\mathbf{e}_1$-adjacent cubes $S, S' \subseteq U \subseteq O_n$ such that (1.20) holds for each of the mappings

$$h_i := \begin{cases} f_1 & \text{if } i = 1, \\ f_i^{-1} \circ f_1 & \text{if } 2 \leq i \leq N, \end{cases}$$

and (1.21) holds for ψ. Using (1.19), we may now write

$$\begin{aligned} \psi(y) = \widetilde{\phi}(y) - \phi(y) &= \left(\widetilde{\phi}(y) - \rho(y)\right) + \rho(y) - \phi(y) \\ &= \left(\widetilde{\phi}(y) - \rho(y)\right) + |\mathrm{Jac}(h_1(y))| - \sum_{i=2}^{N} \rho(h_i(y))\, |\mathrm{Jac}(h_i)(y)| - \phi(y) \end{aligned} \tag{1.24}$$

for a.e. $y \in O_n = Y_1$. To complete the proof we will show that the average value of the final expression over the cube S is too close to its average value over S', that is, closer than the condition (1.21) on ψ allows.

Let $i \in \{2, 3, \ldots, N\}$. Then we have that

$$\|h_i(z) - h_i(y)\|_2 \leq \frac{L}{b(C)} \|z - y\|_2 \leq \frac{L}{b(C)} \operatorname{diam}(U) \leq \delta$$

whenever $y, z \in S \cup S' \subseteq U$. Therefore, in the light of (1.23) and the fact that $|\rho(x) - \phi(x)| = \left|\rho(x) - \widetilde{\phi}(x)\right| \leq \zeta$ for all points x in the image of h_i, we may fix $y_0 \in S$ such that

$$|\rho(h_i(y)) - \rho(h_i(y_0))| \leq |\phi(h_i(y)) - \phi(h_i(y_0))| + 2\zeta \leq 3\zeta \tag{1.25}$$

for all $y \in S \cup S'$. Thus, we have

$$\begin{aligned}
&\left|\fint_S \rho(h_i(y)) \,|\mathrm{Jac}(h_i)(y)| - \fint_{S'} \rho(h_i(y)) \,|\mathrm{Jac}(h_i)(y)|\right| \\
&\leq |\rho(h_i(y_0))| \left|\fint_S |\mathrm{Jac}(h_i)(y)| - \fint_{S'} |\mathrm{Jac}(h_i)(y)|\right| \\
&+ \fint_S \Big(|\rho(h_i(y)) - \rho(h_i(y_0))| \,|\mathrm{Jac}(h_i)(y)|\Big) \\
&+ \fint_{S'} \Big(|\rho(h_i(y)) - \rho(h_i(y_0))| \,|\mathrm{Jac}(h_i)(y)|\Big) \leq \zeta \, \|\rho \circ h_i\|_\infty + 2 \cdot 3\zeta \left(\frac{L}{b(C)}\right)^d .
\end{aligned} \tag{1.26}$$

To derive the final inequality, we bounded the preceding sum term by term using (1.20) for h_i, inequality (1.25) and $|\mathrm{Jac}(h_i)| \leq \left(\frac{L}{b(C)}\right)^d$. The upper bound on the absolute value of the Jacobian of a bilipschitz mappings follows directly from the change of variables for bilipschitz mappings (see Theorem 1.11) and the upper bound on the measure of image of a Lipschitz mapping expressed in equation (1.3) (see 'Background and notation').

For $i \in \{2, \ldots, N\}$, since $\mathrm{supp}(\psi) \subseteq U$ and U is disjoint from $h_i(U)$, we see that $\|\rho \circ h_i - \phi \circ h_i\|_\infty \leq \zeta < 1$, and thus, $\|\rho \circ h_i\|_\infty \leq 1 + \|\phi\|_\infty$. Coming back to equation (1.24) we can combine the last inequality with (1.26) and the fact that $\rho \in B\left(\widetilde{\phi}, \zeta\right)$ to bound the difference of the average values of ψ over S and S' above as

$$\begin{aligned}
\left|\fint_S \psi - \fint_{S'} \psi\right| &\leq 2\zeta + \left|\fint_S |\mathrm{Jac}(h_1)| - \fint_{S'} |\mathrm{Jac}(h_1)|\right| \\
&\quad + (N-1)\zeta \left(\|\phi\|_\infty + 1 + 6\left(\frac{L}{b(C)}\right)^d\right) + \left|\fint_S \phi - \fint_{S'} \phi\right| .
\end{aligned}$$

Note that $\mathrm{diam}(S \cup S') \leq \mathrm{diam}(U) \leq \delta b(C)/L$. Since $b(C) < 1$ (see Theorem 1.35 and Lemma 1.34), we see that $\mathrm{diam}(S \cup S') < \delta$. Using (1.20) for h_1, the fact that $N \leq C$ and the uniform continuity of ϕ in the form of (1.23) we deduce

$$\left|\fint_S \psi - \fint_{S'} \psi\right| \leq 2\zeta + \zeta + C\zeta \left(\|\phi\|_\infty + 1 + 6\left(\frac{L}{b(C)}\right)^d\right) + \zeta. \tag{1.27}$$

However, setting

$$\zeta := \frac{1}{2} \cdot \frac{\varepsilon}{4 + C\left(\|\phi\|_\infty + 1 + 6\left(\frac{L}{b(C)}\right)^d\right)}$$

the right hand side of (1.27) becomes strictly less than ε, contrary to (1.21). Note that ζ is a constant multiple of ε. Thus, we conclude that

$$B\left(\widetilde{\phi}, \zeta\right) \subset B(\phi, 2\varepsilon) \setminus \mathcal{E}_{C,L,n},$$

which demonstrates the porosity of $\mathcal{E}_{C,L,n}$ at ϕ. $\square$

It is now a simple task to combine the previous Lemmas for a proof of Theorem 1.46.

Proof of Theorem 1.46. From Lemma 1.47 we have

$$\mathcal{E} = \bigcup_{C,L\in\mathbb{N}} \mathcal{E}_{C,L} \subseteq \bigcup_{C,L,n\in\mathbb{N}} \mathcal{E}_{C,L,n},$$

whilst Lemma 1.49 asserts that each of the sets in the union on the right hand side is porous. $\square$

Realisability in L^∞ spaces

In the rest of the present section we prove a result similar to Theorem 1.46 about the space $L^\infty(I^d)$, though in a weaker version.

Until now we have only studied realisability in spaces of continuous functions. However, functions ρ admitting a bilipschitz or Lipschitz regular solution $f\colon I^d \to \mathbb{R}^d$ of equation (1.18) need not be continuous. Therefore, it is natural to study the set of realisable functions in the less restrictive setting of $L^\infty(I^d)$, the space of all Lebesgue measurable, real-valued functions ρ defined on I^d, which are essentially bounded. We will prove that the set of all *bilipschitz* realisable functions in $L^\infty(I^d)$ is a σ-porous set. For bilipschitz mappings f, (1.18) is equivalent to the equation

$$|\mathrm{Jac}(f)| = \rho \quad \text{a.e.} \tag{1.28}$$

The question of whether Lipschitz regular realisable densities are also σ-porous, or in some sense negligible, in L^∞ spaces remains open.

Theorem 1.50. *Let*

$$\mathcal{G} := \left\{\rho \in L^\infty(I^d) \colon (1.28) \text{ admits a bilipschitz solution } f\colon I^d \to \mathbb{R}^d\right\}.$$

Then $\mathcal{G}$ is a σ-porous subset of $L^\infty(I^d)$. In fact, $\mathcal{G}$ may be decomposed as a countable union of sets $(\mathcal{G}_L)_{L=1}^\infty$ so that each $\mathcal{G}_L$ is porous at every point $\rho \in L^\infty(I^d)$.

Remark. For $1 \leq p < \infty$, the question of whether the set of bilipschitz realisable densities is small in $L^p(I^d)$ is not interesting because Jacobian of a bilipschitz mapping has to be bounded a.e., as follows from the change of variables for bilipschitz mappings (see Theorem 1.11) and equation (1.3), but the set of all a.e. bounded functions is already σ-porous in the space $L^p(I^d)$.

The proof of Theorem 1.50 will require the following lemma, for which we recall the notation of Section 1.3 (see the statement of Lemma 1.40 and the paragraph above it). The proof itself is a slightly more delicate version of the construction used in the proof of Lemma 1.48.

Lemma 1.51. *Let $\lambda > 0$, $\mathcal{S} \subseteq \mathcal{Q}_\lambda^d$ be a finite collection of tiled cubes in I^d, $\rho \in L^\infty(I^d)$ and $\varepsilon > 0$. Then there exists a function $\psi = \psi(\mathcal{S}, \rho, \varepsilon) \in L^\infty(I^d)$ such that $\|\psi - \rho\|_\infty \leq \varepsilon$ and $\left|⨍_S \psi - ⨍_{S'} \psi\right| \geq \varepsilon$ whenever $S, S' \in \mathcal{S}$ are $\mathbf{e}_1$-adjacent cubes.*

Proof. We define the function ψ on I^d inductively as follows. Pick any $S_1 \in \mathcal{S}$ such that the first co-ordinate projection map $\pi_1 \colon \bigcup \mathcal{S} \to \mathbb{R}$ attains its minimum on S_1. This ensures that $S_1 \neq T + \lambda \mathbf{e}_1$ for any cube $T \in \mathcal{S}$. We set $\psi := \rho$ on S_1.

If $n \geq 1$ and we have already defined ψ on distinct cubes $S_1, \ldots, S_n \in \mathcal{S}$, we extend ψ as follows: if $\mathcal{S} \setminus \{S_1, \ldots, S_n\} = \emptyset$, we complete the construction of ψ by setting $\psi := \rho$ on $I^d \setminus \bigcup \mathcal{S}$. Otherwise, we choose, if possible, $S_{n+1} \in \mathcal{S} \setminus \{S_1, \ldots, S_n\}$ such that $S_{n+1} = S_n + \lambda \mathbf{e}_1$ and define ψ on S_{n+1} by

$$\psi := \begin{cases} \rho & \text{if } \left| \fint_{S_{n+1}} \rho - \fint_{S_n} \psi \right| \geq \varepsilon, \\ \rho + \varepsilon & \text{if } \fint_{S_{n+1}} \rho - \fint_{S_n} \psi \in (0, \varepsilon), \\ \rho - \varepsilon & \text{if } \fint_{S_{n+1}} \rho - \fint_{S_n} \psi \in (-\varepsilon, 0). \end{cases}$$

If it is not possible to find $S_{n+1} \in \mathcal{S} \setminus \{S_1, \ldots, S_n\}$ such that $S_{n+1} = S_n + \lambda \mathbf{e}_1$, we choose $S_{n+1} \in \mathcal{S} \setminus \{S_1, \ldots, S_n\}$ arbitrarily such that the first co-ordinate projection $\pi_1 \colon \bigcup \mathcal{S} \setminus \bigcup_{i=1}^{n} S_i \to \mathbb{R}$ attains its minimum on S_{n+1} and simply take $\psi := \rho$ on S_{n+1}.

It is now readily verified that the function $\psi \in L^\infty(I^d)$ produced by this construction possesses all of the required properties. $\square$

Proof of Theorem 1.50. We decompose $\mathcal{G}$ as $\mathcal{G} = \bigcup_{L=1}^{\infty} \mathcal{G}_L$ where

$$\mathcal{G}_L := \left\{ \rho \in L^\infty(I^d) \colon (1.28) \text{ admits an } L\text{-bilipschitz solution } f \colon I^d \to \mathbb{R}^d \right\}.$$

Fix $L \geq 1$, $\rho \in L^\infty(I^d)$ and $\varepsilon > 0$. We will find $\widetilde{\rho} \in L^\infty(I^d)$ with $\|\widetilde{\rho} - \rho\|_\infty \leq \varepsilon$ and $B(\widetilde{\rho}, \varepsilon/16) \cap \mathcal{G}_L = \emptyset$. This will verify the porosity of the set $\mathcal{G}_L$ and complete the proof of the theorem.

Let $U \subseteq I^d$ be a non-empty, open set, $\zeta = \varepsilon/2$ and let $\eta \in (0,1)$ be a parameter to be determined later in the proof. Let $r \in \mathbb{N}$ and the tiled families $\mathcal{S}_1, \ldots, \mathcal{S}_r$ of cubes contained in U be given by the conclusion of Lemma 1.40 applied to d, $k = 1$, L, η and ζ. We define a sequence of functions $(\widetilde{\rho}_i)_{i=1}^r$ in $L^\infty(I^d)$ by

$$\widetilde{\rho}_i = \psi(\mathcal{S}_i, \rho, \varepsilon) \qquad \text{for } i \in [r],$$

where $\psi(\mathcal{S}_i, \rho, \varepsilon)$ is given by the conclusion of Lemma 1.51. Now let $\widetilde{\rho} \in L^\infty(I^d)$ be defined by

$$\widetilde{\rho}(x) = \begin{cases} \widetilde{\rho}_i(x) & \text{if } x \in \bigcup \mathcal{S}_i \setminus \bigcup_{j=i+1}^{r} \bigcup \mathcal{S}_j,\ i \in [r], \\ \rho(x) & \text{if } x \in I^d \setminus \bigcup_{i=1}^{r} \bigcup \mathcal{S}_i. \end{cases}$$

It is clear that $\|\widetilde{\rho} - \rho\|_\infty \leq \varepsilon$. Let $\phi \in B(\widetilde{\rho}, \varepsilon/16)$. Then, given $i \in [r]$ and $\mathbf{e}_1$-adjacent cubes $S, S' \in \mathcal{S}_i$, we let $T := S \cap \bigcup_{j=i+1}^{r} \bigcup \mathcal{S}_j$ and $T' := S' \cap \bigcup_{j=i+1}^{r} \bigcup \mathcal{S}_j$. From Lemma 1.40, part 1 we have that $\max\{\mathcal{L}(T), \mathcal{L}(T')\} \leq \eta \mathcal{L}(S)$. We deduce

$$\begin{aligned} \left| \fint_S \phi - \fint_{S'} \phi \right| \geq & \left| \fint_S \widetilde{\rho}_i - \fint_{S'} \widetilde{\rho}_i \right| - \left| \fint_S (\phi - \widetilde{\rho}) - \fint_{S'} (\phi - \widetilde{\rho}) \right| \\ & - \frac{1}{\mathcal{L}(S)} \left| \int_T (\widetilde{\rho} - \widetilde{\rho}_i) - \int_{T'} (\widetilde{\rho} - \widetilde{\rho}_i) \right| \geq \varepsilon - \frac{2\varepsilon}{16} - 4\varepsilon\eta > \frac{\varepsilon}{2}, \end{aligned}$$

when we set $\eta = \frac{1}{16}$. Together with Lemma 1.40, part 2 and setting $\zeta = \varepsilon/2$, this implies that equation (1.28) with $\rho = \phi$ has no L-bilipschitz solutions $f \colon I^d \to \mathbb{R}^d$. Hence $\phi \notin \mathcal{G}_L$. $\square$

Conclusion

To conclude the chapter, we summarise the most important results obtained so far, comment on possible further directions for research and present several open questions.

The main result presented in this chapter is the negative answer to Feige's question (Question 1.4) in any dimension $d \geq 2$; indeed, a generic positive, continuous function on I^d which is also bounded away from zero provides a counter-example when used in the modified discretisation procedure of Burago and Kleiner [8] presented in Section 1.1. However, this procedure is not restricted only to continuous functions. As was shown in Theorem 1.16, any bounded, measurable function $\rho\colon I^d \to [0, \infty)$ that is positive a.e. can be discretised. And if there is no Lipschitz mapping $f\colon I^d \to \mathbb{R}^d$ satisfying $f_\sharp \rho \mathcal{L} = \mathcal{L}|_{f(I^d)}$, then ρ provides a counter-example to Feige's question.

In the spirit of the work of Burago and Kleiner [8] and McMullen [42], in Question 1.5 we have formulated a continuous analogue of Question 1.4. However, we were able to establish only a one-sided relation between the two questions. Namely, in Theorem 1.6 we have shown that a negative answer to the continuous question implies a negative answer to the discrete question. So it is natural to ask whether there is a relation in the opposite direction as well. The author would like to describe informally some of his thoughts related to it. The following discussion is mathematically a bit vague, but the ideas will hopefully be clear.

All counter-examples produced via Theorem 1.6 are of a rather special form. Discrete sets $S \subset \mathbb{Z}^d$ of size $|S| = n^d$ obtained this way are 'almost cubical'; that is, they are contained in a cubical grid of side $O(n)$, where the implicit constant hidden in $O(n)$ depends only on $\sup \rho$ and d. On the one hand, it is a remarkable fact that it is possible to produce sequences $(S_n)_{n\in\mathbb{N}}$ providing a counter-example to Feige's question such that S_n fits in a cubical grid of side at most cn for any $c > 1$ arbitrarily close to 1. On the other hand, the special property of sets produced this way suggests that Question 1.5 may *not* be equivalent to Feige's question; Question 1.15 allows for much greater variety of sets than just 'almost cubical' ones.

One could hope for a procedure that takes any sequence of discrete sets $S_n \subset \mathbb{Z}^d$ with $|S_n|$ tending to infinity with n and encodes it into a measurable function $\rho\colon I^d \to [0, \infty)$. Then one would like to establish a relation between existence of Lipschitz solutions $f\colon I^d \to \mathbb{R}^d$ to the equation $f_\sharp \rho \mathcal{L} = \mathcal{L}|_{f(I^d)}$ on the one side and existence of Lipschitz bijections $S_n \to [a_n]^d$ with a bounded Lipschitz constant on the other side, where $a_n^d = |S_n|$.

The author has tried to come up with such a procedure, but failed. A natural approach seems to be to take a sequence of sets $S_n \subset \mathbb{Z}^d$ and try to encode S_n into a measure with support in I^d. For instance, one can scale down S_n so that it fits inside I^d and then assign a certain weight to each point, or take a *Voronoi diagram* generated by S_n and use its cells to define a measure. Using the weak compactness of the space of compactly supported probability measures, one can then extract a subsequence of these measures and look at their weak limit. However, here comes the trouble: a sequence of sets $(S_n)_{n\in\mathbb{N}}$ that is not 'almost cubical' in the sense explained above usually yields a measure that is *singular* with respect to the Lebesgue measure. Thus, it does not have any density ρ with respect to the Lebesgue measure and one does not end up with a function to work with.

One possible way to circumvent this trouble could be to try to come up with a *kernelization* argument:

Question 1.52 (Open). *Is it true that sequences of 'almost cubical' sets $S_n \subset \mathbb{Z}^d$ provide, in some sense, the hardest instances for Feige's question? For instance, is there $L > 0$ such that for every sequence of sets $S_n \subset \mathbb{Z}^d$, $|S_n| = a_n^d, a_n \in \mathbb{N}$ such that a_n goes to infinity with n and S_n fits only inside cubical grids of side $\omega(a_n)$ there exists a sequence of L-Lipschitz bijections $S_n \to [a_n]^d$? In other words, is the answer to Feige's question positive for sets that are not 'almost cubical'?*

Since we already know that the answer to Feige's question is negative, it would be interesting to know what is the worst possible rate of growth of minimal Lipschitz constants for counter-examples to Feige's question.

Question 1.53 (Open). *Let*

$$L_n := \min\left\{L > 0 : \left(\forall S \subset \mathbb{Z}^d, |S| = n^d\right)\left(\exists L\text{-Lipschitz bijection } S \to [n]^d\right)\right\}.$$

What is the rate of growth of $(L_n)_{n=1}^{\infty}$?

Feige's approximation algorithm for the Graph Bandwidth, which was the original motivation for Question 1.4, provides $O(\text{polylog}(n))$ approximation ratio, where n is the number of vertices. Thus, it would be interesting to know whether $L_n \in O(\text{polylog}(n))$ or not[18]. Currently, only a trivial bound $L_n \in O(n)$ is known.

Theorem 1.46 asserts that a typical continuous function $I^d \to (0, \infty)$ is not Lipschitz regular realisable. This is a twofold generalisation of the result of Burago and Kleiner [8], who proved the existence of continuous bilipschitz non-realisable function.

However, Theorem 1.46 is only an existence result; it does not provide a concrete example of Lipschitz regular non-realisable function. Curiously, the author does not know how to obtain such an example.

It may also be a bit surprising that in the less restrictive setting of $L^\infty(I^d)$ we were able to prove typicality only for bilipschitz non-realisable functions instead of Lipschitz regular non-realisable functions.

The main reason lies in the following. Recall that, by Theorem 1.35, a Lipschitz regular mapping $f : I^d \to \mathbb{R}^d$ decomposes on some open subset of $f(I^d)$ precisely into a sum of bilipschitz homeomorphisms $f_1, \ldots, f_N$. Thus, assuming that O_n is the domain of f_1, the equation $f_\sharp \rho \mathcal{L} = \mathcal{L}|_{f(I^d)}$ yields the following relation between ρ and Jacobians of f_i, as was shown in the proof of Lemma 1.47 in Section 1.4:

$$\rho(y) = |\text{Jac}(f_1)(y)| - \sum_{i=2}^{N} \rho\left(f_i^{-1} \circ f_1(y)\right)\left|\text{Jac}(f_i^{-1} \circ f_1)(y)\right| \qquad \text{for a.e. } y \in O_n. \tag{1.19 revisited}$$

[18]We leave aside the question of finding a specific Lipschitz bijection with Lipschitz constant $O(L_n)$ effectively; it is not clear at all how to do this.

For a bilipschitz mapping f the whole sum in the equation above simply disappears and we can easily control the remaining term using Lemma 1.40, statement 2. While for a Lipschitz regular mapping f Lemma 1.40 provides some control over $\left|\mathrm{Jac}(f_i^{-1} \circ f_1)\right|$, we need to work with $\left(\rho \circ f_i^{-1} \circ f_1\right) \left|\mathrm{Jac}(f_i^{-1} \circ f_1)\right|$ instead. In the case of continuous ρ it can be remedied using the uniform continuity of ρ; prescribing sufficiently small O_n at the beginning, we can treat ρ in equation (1.19) essentially as a constant. However, in the case of $\rho \in L^\infty(I^d)$ we have no local control over the behaviour of ρ, and thus, we are not able to use equation (1.19) in conjunction with Lemma 1.40 to prove that a typical essentially bounded measurable function on I^d is not Lipschitz regular realisable.

Question 1.54 (Open)**.** *What is the size of the set of Lipschitz regular non-realisable functions inside the space $L^\infty(I^d)$? Is it residual*[19]*? Is it σ-porous?*

[19] That is, of second category.

2. The Hanani–Tutte theorem on the projective plane

Introduction

The study of the most diverse ways how to draw a graph has been a classical and increasingly popular part of mathematics at the intersection of discrete mathematics, geometry, topology as well as computational complexity. The present chapter falls into the area of drawing of graphs on surfaces.

History and motivation

Following the work of Chojnacki [5] from 1934, who later changed his name to Hanani, Tutte [30] made a remarkable observation in 1970, which became known as the strong Hanani–Tutte theorem:

Theorem 2.1 (Strong Hanani–Tutte theorem [5, 30])**.** *A graph is planar if and only if it can be drawn*[1] *in the plane in such a way that every two non-adjacent*[2] *edges cross an even number of times, possibly zero times.*

Of course, the 'only if' part in the theorem above is trivial. The reason why the theorem is called 'strong' is that there also exists a 'weak' version: a graph is planar if and only if it has a drawing in the plane in which every two edges cross an even number of times.

It is natural to ask whether the strong Hanani–Tutte theorem can be extended to drawings of graphs on *surfaces* different from the plane.

Question 2.2 (Strong Hanani–Tutte for surfaces)**.** *Given a closed surface S, is it true that a graph G is embeddable into S if and only if it has a drawing on S in which every two non-adjacent edges cross an even number of times?*

Although it seems plausible that the question has been known to the graph drawing community for a few decades, to the best of author's knowledge, it appeared explicitly only in [28, Conj. 1], where Schaefer and Štefankovič conjectured that the answer was positive.

Currently, the positive answer to Question 2.2 is only known for the sphere (plane) and for the *projective plane*. On the other hand, in 2017 Fulek and Kynčl [9] have provided a negative answer for all *orientable* surfaces of *genus* four and higher. The remaining cases, which comprise of the torus, the double-torus and the triple-torus as well as of all non-orientable surfaces of genus two and higher, are open.

In contrast, for the weak version of the Hanani–Tutte theorem the situation is quite different. In 2000 Cairns and Nikolayevsky [4] proved that the weak version holds for all orientable surfaces; later, in 2007, Pelsmajer, Schaefer, and Štefankovič [23] established the result for all surfaces. It is worth noting that the weak version of the theorem allows for stronger conclusion than the strong version:

[1] We consider only drawings in which distinct vertices are represented by distinct points and each edge is represented by an arc which does not go through any vertex except at the endpoints. Moreover, two such arcs are allowed to cross only finitely many times.

[2] Two edges are called *non-adjacent* if they do not share a vertex.

Theorem 2.3 (Weak Hanani–Tutte theorem on surfaces [4, 23])**.** *If a graph is drawn on a surface in such a way that every pair of edges crosses an even number of times, then the graph can be embedded into that surface while preserving the* cyclic order *of the edges at all vertices*[3]*.*

Note that in the strong version we require that only non-adjacent edges cross even number of times, while in the weak version this condition has to hold for all pairs of edges. Consequently, the strong version and the weak version of the theorem are not comparable.

The original proof of Theorem 2.1 by Hanani [5] and Tutte [30] uses Kuratowski's theorem [16], that is, a characterisation of planar graphs as those graphs that do not contain the graphs K_5 and $K_{3,3}$ as minors. The strong Hanani–Tutte theorem has also a parallel history in algebraic topology in works of van Kampen [31], Wu [32], Shapiro [29] and Levow [17]. For a complete history and relevant results we refer to a nice survey by Schaefer [26].

In 2007, Pelsmajer, Schaefer, and Štefankovič [22] provided the first constructive proof of the strong Hanani–Tutte theorem. They presented a procedure that takes a drawing of a graph in the plane such that non-adjacent edges cross even number of times and redraws it gradually into an embedding.

In 2009 Pelsmajer, Schaefer, and Stasi [24] proved a version of the strong Hanani–Tutte theorem for the projective plane:

Theorem 2.4 (Strong Hanani–Tutte for the projective plane [24])**.** *A graph can be embedded into the projective plane if and only if it can be drawn on the projective plane in such a way that every two non-adjacent edges cross an even number of times.*

The proof of Theorem 2.4 by Pelsmajer et al. [24] is based on a characterisation of graphs embeddable into the projective plane via minimal forbidden minors. This approach is relatively simple for the projective plane; however, it does not seem applicable to other surfaces. On the one hand, in their famous work, Robertson and Seymour [25] have proved that for every surface S there is a finite list $F(S)$ of graphs such that a graph G is embeddable into S if and only if none of the minors of G is on the list $F(S)$. On the other hand, there is no reasonable characterisation of forbidden minors for surfaces other than the plane and the projective plane. The complete list of the minimal forbidden minors for the projective plane consists of 35 graphs found by Glover, Huneke, and Wang [13] and Archdeacon [2]. Already for the torus or the Klein bottle the exact lists are not known, but it is known that the list for the torus contains thousands of graphs (see the work of Gagarin, Myrvold, and Chambers [12] and the references therein). Moreover, by additivity of the genus of a graph established by Battle, Harary, Kodama, and Youngs [3] and Miller [18], the number of minimal forbidden minors for a surface is at least exponential in its genus.

The results in the present chapter come from the work of Colin de Verdière, the author, Paták, Patáková, and Tancer [6]; the main result is a *constructive* proof of Theorem 2.4. The need for such a proof is motivated by the unsolved[4]

[3]In fact, the embedding preserves the so-called *embedding scheme* of the given drawing of the graph on the surface. For the precise definitions, see, e.g., Mohar and Thomassen [20].

[4]The work of Fulek and Kynčl [9] was published after the first version of the present work was finished; at that time, the only solved cases were the plane and the projective plane.

cases of Question 2.2.

For simplicity, we call a drawing of a graph on a surface a *Hanani–Tutte drawing* if no two non-adjacent edges cross an odd number of times[5].

Given a Hanani–Tutte drawing of a graph on the projective plane, the proof of Colin de Verdière, the author, Paták, Patáková, and Tancer [6], which we present here, gives an explicit way to transform the drawing into an embedding. In principle, the proof could be transformed into a polynomial-time algorithm. On the other hand, there already are linear-time algorithms for deciding embeddability of a graph on the projective plane by Mohar [19] and by Kawarabayashi, Mohar, and Reed [15] (these algorithms work for any surface, but the hidden constant depends exponentially on the genus).

The presented approach reveals a number of difficulties that have to be overcome in order to obtain a constructive proof. If the answer to Question 2.2 is affirmative for some surface S, the ideas of Colin de Verdière, the author, Paták, Patáková, and Tancer [6] may serve as a basis for the proof of this fact. If it is negative, then an approach based on these ideas may help reveal an appropriate structure needed to construct a counter-example.

We remark that other variants of the Hanani–Tutte theorem generalising the notion of embedding in the plane have also been considered. For instance, Schaefer [27] established a version of the strong Hanani–Tutte theorem for *partially embedded graphs* and versions of both the weak and the strong Hanani–Tutte theorem were also proven for 2-*clustered* graphs by Fulek, Kynčl, Malinović, and Pálvölgyi [11]. For more details and other variants we again refer to the survey by Schaefer [26].

One of the reasons why the strong Hanani–Tutte theorem is important is that it turns the question of planarity of a given graph into a system of linear equations. For general surfaces, the existence of a Hanani–Tutte drawing of G leads to a system of quadratic equations over $\mathbb{Z}_2$ (see Levow [17]). If the strong Hanani–Tutte theorem is true for the surface, any solution to the system then serves as a certificate that G is embeddable into that surface. Moreover, if the proof of the Hanani–Tutte theorem is constructive, it gives a recipe how to turn the solution into an actual embedding. Unfortunately, solving systems of quadratic equations over $\mathbb{Z}_2$ is **NP**-complete in general.

Redrawing procedure

Let us call an edge *even* in a drawing of a graph if it crosses every other edge an even number of times (including the adjacent edges).

As we noted before, Pelsmajer et al. [22] provided the first constructive proof of the strong Hanani–Tutte theorem in the plane. They showed a sequence of moves transforming a Hanani–Tutte drawing into an embedding. A key step in their proof is the following theorem:

Theorem 2.5 (Pelsmajer et al. [22, Thm. 2.1]). *If D is a drawing of a graph G in the plane and E_0 is the set of even edges in D, then G can be drawn in the plane so that no edge in E_0 is involved in an intersection and there are no new pairs of edges that intersect an odd number of times.*

[5] Such a drawing is also called an independently even drawing in the literature.

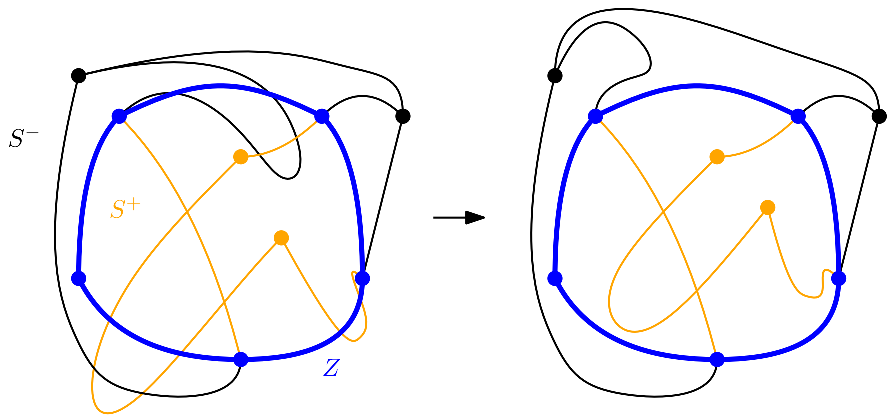

Figure 2.1: Separating the outside (in black) and the inside (in orange) of a cycle Z (in blue; thick). The symbols S^+ and S^- denote the two parts of the plane separated by Z.

Unfortunately, a result analogous to Theorem 2.5 is simply not true on other surfaces, as was shown by Pelsmajer et al. [23]. In particular, this is an obstacle for a constructive proof of the strong Hanani–Tutte theorem for the projective plane (Theorem 2.4).

The key step of the approach presented here is to provide a suitable replacement of Theorem 2.5 ([22, Thm. 2.1]; see also Lemma 3 in [10]).

To explain the basic idea of the suggested replacement, let us first describe the situation on the sphere. Consider a graph G with a Hanani–Tutte drawing D on the sphere. Let Z be a cycle in G which is drawn without self-intersections and such that every edge of Z is even. Theorem 2.5 then implies that G can be redrawn in such a way that Z becomes free of crossings in the new drawing without introducing any new pairs of edges crossing oddly.

In fact, a detailed inspection of the redrawing procedure from Theorem 2.5 reveals something slightly stronger in this setting. The drawing of Z splits the plane (or the sphere) into two parts that we call the *inside* and the *outside*. This in turn splits G into two parts. The inside part consists of vertices that are inside Z and of the edges that have either at least one endpoint inside Z, or they have both endpoints on Z and point to the inside of Z from both endpoints. The outside part is defined analogously. Because we have started with a Hanani–Tutte drawing, it is easy to check that every vertex and every edge is on Z, inside, or outside. The proof of Theorem 2.5 in [22] then implies that it is possible to fully separate the inside and the outside in the drawing; see Figure 2.1. Actually, the separation can be performed even by a continuous motion of each of the parts, provided the drawing is considered on the sphere instead of the plane.

The trouble on the projective plane is that it may not be possible to separate the outside and the inside of a separating cycle by a continuous motion (of each of the parts separately). This is demonstrated by a projective-planar drawing of K_5 in Figure 2.2, left. The symbol '$\otimes$' stands for the *crosscap*[6] in the picture.

[6]We can think of a crosscap as a small disk whose interior is removed and the opposite points on its boundary are identified. The projective plane can be thought of as a sphere with one crosscap.

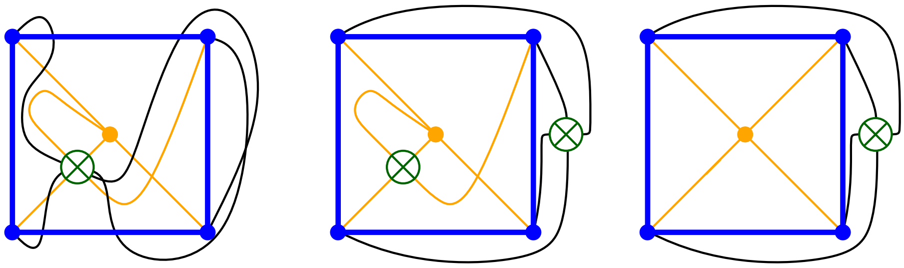

Figure 2.2: A drawing of K_5 on the projective plane in which the outside and the inside cannot be separated by a continuous motion (left). A solution by duplicating the crosscap (middle) and removing one of them in the next step (right).

We could easily move the part of the graph to the outside as desired if we were allowed to duplicate the crosscap as in Figure 2.2, middle. However, the problem is that the sphere with two crosscaps is not the projective plane, but the Klein bottle, that is, a different surface. On the other hand, if we give up on a continuous motion, we may observe that the inside vertices and edges in Figure 2.2, middle, may actually be redrawn without using a crosscap at all; we can then simply remove the 'inside' crosscap. This step changes the *homotopy/homology type* of many cycles in the drawing.

The main technical contribution of the present work is to show that it was not a coincidence that a simplification as in the drawing in Figure 2.2 was possible. It will be shown that we can always redraw one of the sides without using the 'duplicated' crosscap. The precise statement requires some preparation and will be given later in Theorem 2.14 in Section 2.2.

The overall idea of the proof of the strong Hanani–Tutte theorem in the plane of Pelsmajer et al. [22] is to find a suitable order on some of the cycles of the graph and then use Theorem 2.5 repeatedly on these cycles, which eventually results in an embedding in the plane. A detailed proof of Pelsmajer et al. [22] uses an induction based on this idea.

Similarly, the presented proof of Theorem 2.4 uses inductively Theorem 2.14 on suitably chosen cycles. However, the details are more complicated compared to the setting of [22], because on the projective plane it is necessary to deal with two types of cycles in the graph based on their homology. It is also more delicate to set up the induction in a way suitable for Theorem 2.14, since the setting of Theorem 2.14 is slightly more restrictive than the setting of Theorem 2.5.

Organisation of the chapter

To present the approach of Colin de Verdière, the author, Paták, Patáková, and Tancer [6] we need to develop an appropriate toolbox for manipulation with Hanani–Tutte drawings on the projective plane. Actually, many of the tools are applicable to a general surface. In order to motivate the introduction of all the auxiliary definitions and lemmas, the main proofs are presented as soon as possible, which is reflected in the structure of the chapter.

Therefore, in Section 2.1 only basic definitions and facts are presented, which include tools to modify drawings and to represent the Hanani–Tutte drawings on the projective plane as drawings on the sphere.

In Section 2.2 the precise statement of Theorem 2.14, the separation theorem, is described together with its proof. However, the proofs of many auxiliary results are postponed to later sections.

Section 2.3 is devoted to the proof of Theorem 2.4, the strong Hanani–Tutte theorem for the projective plane, using Theorem 2.14 and some of the auxiliary results from Section 2.2.

Section 2.4 contains a description of additional technical tools needed for the proofs of the auxiliary results, which are then presented gradually in the remaining sections.

A general background, facts and notation related to graphs, surfaces, drawings and basic homology, which will be used throughout the chapter, are presented in 'Background and notation' following this one.

Background and notation

In this section we recall basic terms, notation and facts used in graph theory and theory of surfaces as well as a few basic facts from homology. The terminology and notation specific to the present work will be introduced gradually in the subsequent sections.

For the readers not familiar with the topic, the author would like to remark that it is often sufficient to rely on an intuitive understanding of the terms used here (like a surface, a drawing or an embedding of a graph).

Graphs

As a general reference for graph theory we suggest a well-known monograph by Diestel [7].

Let G be a graph. We denote by $V(G)$ the set of vertices of G and by $E(G)$ the set of edges of G. We consider only *simple, undirected* graphs, that is, G can have no loops nor parallel edges and every edge is an unordered pair of vertices. To simplify the notation, for an edge e with endpoints u, v we usually write $e = uv$ instead of $e = \{u, v\}$.

We say that two edges are *independent* if they do not share a vertex. We say that an edge e is *incident* to a vertex v if v is one of the endpoints of e.

Given a vertex $v \in V(G)$, we denote by $G - v$ the graph obtained from G by removing v and all edges adjacent to v. Similarly, for $e \in E(G)$ we write $G - e$ for the graph obtained from G by removing the edge e, but preserving its endpoints. More generally, for $W \subset V(G)$ we write $G - W$ for the graph we get by removing all vertices of W from G. Similarly, we use $G - F$ for the graph G without the edges in F, where $F \subseteq E(G)$. For $W \subset V(G)$ we call a subgraph G' of G *induced* by W if $V(G') = W$ and $e \in E(G')$ if and only if $e \in E(G)$ and both endpoints of e are in W.

By a *walk* ω in a graph G we mean a sequence $v_0, e_1, v_1, e_2, v_2, \dots, e_n, v_n$ such that $v_i \in V(G)$ for every $i \in \{0, \dots, n\}$ and such that $e_i \in E(G), e_i = \{v_{i-1}, v_i\}$ for every $i \in \{1, \dots, n\}$. The vertices v_0 and v_n are called the *endpoints* of ω, and sometimes, we call ω a v_0v_n*-walk*. We say that ω is *closed* if $v_0 = v_n$. The length of a walk is defined as the number of edges it traverses, i.e., in the case above the length of ω is n. We also allow for walks of length zero, which consist of a single vertex.

If a walk ω does not visit any edge nor any vertex more than once, except the possibility that it starts and ends in the same vertex, then we call it a *path* in G. A closed path is called a *cycle* in G.

A graph G is *connected* if for every two vertices $u, v \in V(G)$ there is a path in G with endpoints u and v. The inclusion-maximal connected subgraphs of G are called the *connected components* of G. Thus, a connected graph has exactly one connected component. If a connected component consists of a single vertex, the vertex is called *isolated*.

A graph G is (vertex) k*-connected* if $|V(G)| > k$ and $G - W$ is connected for every $W \subset V(G)$ with $|W| < k$. Dually, by *Menger's theorem* (see, e.g., Diestel [7, Cor. 3.3.5]), G is k-connected if every two distinct vertices of G can be connected by k paths which are disjoint up to their endpoints. In fact, Menger's theorem ensures something stronger: for any two sets $A, B \subseteq V(G)$ the minimum number of vertices of G that have to be removed to separate A from B is equal to the maximum number of vertex-disjoint paths between A and B. A set of vertices the removal of which partitions a connected graph into multiple connected components is called a *cut*. If a cut consist of a single vertex, the vertex is called a *cut vertex*.

A graph is called a *forest* if it does not contain any cycles. It is called a *tree* if it is a connected forest. A subgraph T of a graph G is called a *spanning forest* of G or a *spanning tree* of G, if $V(T) = V(G)$ and T is an edge-maximal forest or a tree, respectively.

Surfaces

For a detailed treatment of closed surfaces and drawings of graphs on surfaces we refer the reader to Mohar and Thomassen [20].

We write S^2 for the 2-dimensional sphere. We can *attach a crosscap* to S^2 by choosing a small disk (i.e., 2-dimensional ball) B in S^2, removing its interior and identifying the opposite points on the boundary of B. The *crosscap* is then the curve obtained from the boundary of B after the identification. In figures, we use the symbol '$\otimes$' for the crosscap coming from the removal of the disk 'inside' this symbol. We note that the sphere with a crosscap cannot be embedded in Euclidean space of dimension less than 4.

We can also attach a *handle* to S^2. This is done by choosing two disjoint small disks B_1 and B_2 in S^2, removing their interiors and gluing their boundaries together in a one-to-one way; the direction in which we glue the boundary of B_1 to the boundary of B_2 is the opposite to the direction chosen on the boundary of B_2. The result should look very much like an ordinary handle attached to a cup.

A usual textbook definition of a *closed surface* says that it is a connected, compact, Hausdorff, second countable topological space that is locally homeomorphic to an open disk in $\mathbb{R}^2$. Perhaps a more comprehensible equivalent characterisation of closed surfaces is provided by the *classification of surfaces* (see, e.g., Mohar and Thomassen [20])—every closed surface is homeomorphic either to S^2 with h handles, or to S^2 with k crosscaps for some values of $h, k \in \{0, 1, \ldots\}$. We will consider only closed surfaces, and thus, we refer to them simply as surfaces. The sphere S^2 with one crosscap (and zero handles) attached to it is called the (real) projective plane and denoted by $\mathbb{R}P^2$.

The surfaces created from S^2 by attaching handles and no crosscaps are called *orientable*, while by attaching at least one crosscap we produce a *non-orientable* surface. Thus, the projective plane is non-orientable. On an orientable surface it is possible to assign an orientation to a neighbourhood of every point in a globally consistent way, while in the non-orientable case this is impossible.

The number of handles attached to S^2 to obtain an orientable surface S is called the *orientable genus* of S. Similarly, the number of crosscaps used to produce a non-orientable surface S is called the *non-orientable* genus of S. We define the *genus*[7] of a surface S to be equal to its non-orientable genus if S is non-orientable, and equal to *twice* its orientable genus if it is orientable.

[7] In literature, this is usually called *Euler genus*.

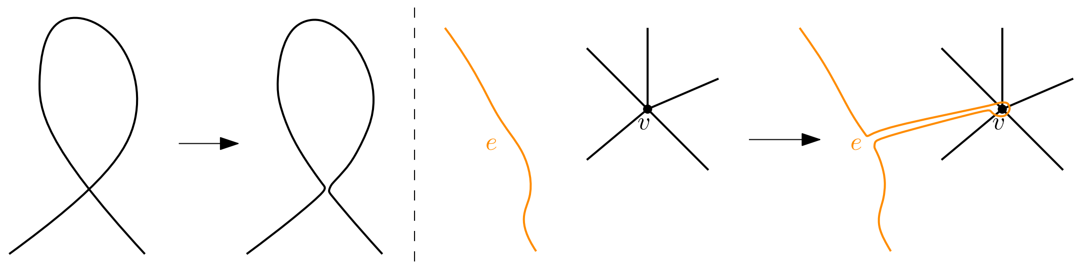

Figure 2.3: (Left) The way to eliminate a self-crossing of an edge. (Right) A vertex-edge switch (v, e).

Drawings of graphs and crossings

Intuitively, a *drawing* of a graph G on a surface S is a representation of G in which the vertices of G are represented by distinct points on S and each edge of G is represented by a curve on S connecting the points representing the endpoints of the edge. An *embedding* of G into S is then a drawing of G in which the curves representing the edges are not allowed to cross nor touch, except at the endpoints.

To provide more formal definitions, we will think of a graph G as of one-dimensional simplicial complex. Then an *embedding* of G on a surface S is an injective continuous map from a geometric realisation of the simplicial complex representing G to S that is a homeomorphism onto its image.

A *drawing* of G on S is any continuous map from the simplicial complex representing G to S that is injective on vertices. However, we will not consider drawings in such a generality. We put the standard general position assumptions on the drawings. That is, we consider only drawings of graphs such that no edge (i.e., the image of the edge in the drawing) contains a vertex in its interior and every pair of edges meets only in a finite number of points, where they intersect transversally. We refer to a transversal intersection of two curves as to a *crossing* or we simply say that the two curves *cross* at the point.

Since we study only pairwise interactions of edges, we allow three or more edges crossing in a single point. Moreover, note that it is possible to get rid of a self-crossing of an edge without changing the image of the edge except in a small neighbourhood of the self-crossing; see Figure 2.3, left. Therefore, throughout the chapter we assume that whenever a self-crossing of an edge appears in a drawing it is immediately eliminated.

A graph that can be embedded into S^2 (or the plane) is called *planar*. A graph embeddable into the projective plane is called *projective-planar*. Let D be an embedding of a graph G into a surface S. The connected components of $S \setminus D$ are called *faces* of the embedding D.

Given a surface S and a graph G, we recall that a *Hanani—Tutte drawing* of G on S is a drawing of G on S such that every pair of independent edges crosses an even number of times. We will often abbreviate the term Hanani–Tutte drawing to *HT-drawing*.

Given an HT-drawing of a graph on $\mathbb{R}P^2$, we can slightly alter it in such a way that every edge meets the crosscap in a finite number of points and only transversally, still keeping the property that we have an HT-drawing. Therefore, we may add to our assumptions that this is the case for all HT-drawings on $\mathbb{R}P^2$

that we shall consider.

To every path γ in a graph G corresponds a curve in a drawing D of G, which we denote by $D(\gamma)$. More generally, for a subgraph H of the graph G we write $D(H)$ for the part of the drawing D corresponding to the graph H. We can extend the notation to walks as well: given a walk $\omega = v_0, e_1, \dots, e_n, v_n$ in G, we denote by $D(\omega)$ the closed curve, possibly self-overlapping, that is defined as the concatenation of $D(e_i)$, for $i = 1, \dots, n$ in the directions determined by ω.

We call a path (or a cycle) γ of G *simple* in D if the curve $D(\gamma)$ does not intersect itself (except possibly at the endpoints in the case of a cycle).

Let D be a drawing of a graph G. Given two distinct edges e and f of G, by $\mathrm{cr}_D(e, f)$ we denote the number of crossings between e and f in D modulo 2. If the drawing D is clear from a context, we usually drop the subscript D and write just $\mathrm{cr}(e, f)$. We say that an edge e of G is *even* if $\mathrm{cr}(e, f) = 0$ for every $f \in E(G)$ distinct from e. We emphasize that we treat the crossing number as an element of $\mathbb{Z}_2$ and all computations throughout the chapter involving it are done in $\mathbb{Z}_2$.

Let D be a drawing of a graph G on S^2. Let us consider a vertex $v \in V(G)$ and an edge $e \in E(G)$ such that v is not incident to e. Imagine we pull a thin 'finger' from the interior of e towards v and we let this finger pass over v. See Figure 2.3, right. Let us write D' for the new drawing. We say that D' is obtained from D by the *vertex-edge switch*[8] (v, e). For any edge f incident to v the crossing number $\mathrm{cr}(e, f)$ changes from 0 to 1, or vice versa. It does not change for the other pairs of edges, because the 'finger' from e intersects the other edges in pairs of points.

Homology of surfaces

The basics of homology of surfaces are covered by Munkres [21]. A more general and detailed treatment of the homology theory can be found in Hatcher [14].

We use singular homology with $\mathbb{Z}_2$ coefficients unless specified otherwise. Nevertheless, the work presented in this chapter should be understandable even to those that are not familiar with any homology or homotopy theory. The properties needed in the present work are summarised bellow. The following discussion is primarily intended as a rather informal explanation for the readers not familiar with the homology theory.

We will use homology only as a classification of closed curves on a surface. The homology on the sphere is trivial, which means that all closed curves on the sphere have the same type in homology. The homology on the projective plane distinguishes between two types of closed curves: those that pass over the crosscap an even number of times and those that pass it an odd number of times. The former are called *trivial*, while the latter are called *non-trivial*. We may extend the definition to any linear combination of curves by linearity—the type of the linear combination will be the same combination of the respective types of curves (we think of the trivial curves having the type 0 and the non-trivial ones having the type 1 and the calculations are done in $\mathbb{Z}_2$).

[8]Another name for the vertex-edge switch is the *finger-move* common mainly in topological context in higher dimensions.

The crosscap is homologically non-trivial simple, closed curve in $\mathbb{R}P^2$. An important fact from the homology theory is that the distinction between trivial and non-trivial curves is independent of the choice of the crosscap used to represent the projective plane. This, in particular, means that any homologically non-trivial simple, closed curve may serve as a crosscap up to a self-homeomorphism of $\mathbb{R}P^2$.

Whenever we speak about the homology type of a cycle in a drawing of a graph, we mean the type of the curve representing the cycle. We can also speak about the homology type of a closed walk, since the concatenation of the curves representing the edges of the walk gives a closed curve on the surface.

We also note that the trivial curves on $\mathbb{R}P^2$ are *contractible* meaning they may be continuously shrunk to a point. Trivial, simple curves are also *separating*, that is, cutting the surface along such a curve yields two connected components. The non-trivial curves are *non-contractible* and simple, non-trivial curves are *non-separating.*

The last general fact from the homology theory that we shall need comes from the so-called *intersection form* on surfaces. For a precise definition we refer the reader to Fuchs and Viro [8, Sect. 8.4]. It is sufficient for us to state here only what it says about the sphere (the plane) and the projective plane.

In the case of the sphere the intersection form is trivial, which means that any two closed curves on the sphere which intersect only finitely many times and cross at every intersection have to cross an even number of times[9] [10].

However, in the case of the projective plane the intersection form is non-trivial as explained in the following lemma:

Lemma 2.6 (Intersection form on $\mathbb{R}P^2$)**.** *Let z_1 and z_2 be two closed curves on $\mathbb{R}P^2$ that intersect only finitely many times and cross at every intersection. Then z_1 and z_2 cross an odd number of times if and only if they are both homologically non-trivial.*

Many of the considerations present in the subsequent sections rely heavily on Lemma 2.6.

[9]We emphasize that we think of closed curves as of continuous mappings from the unit circle to the surface. Thus, we count their intersections as the number of pairs of points from their domains that are mapped to the same point on the surface. This means that we count the common points in their images *with multiplicity.*

[10]This is also a consequence of the Jordan curve theorem (see, e.g., Mohar and Thomassen [20, p. 25]).

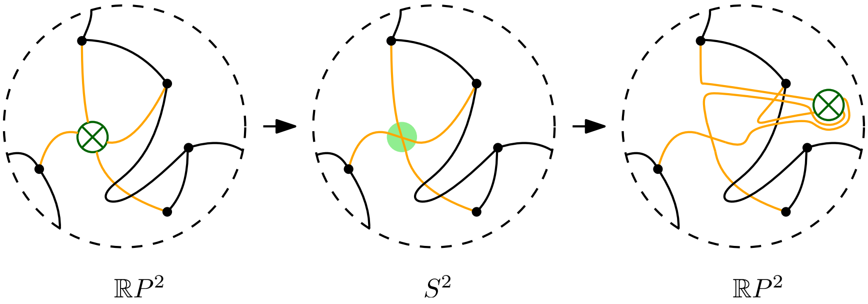

Figure 2.4: A transformation between an HT-drawing on $\mathbb{R}P^2$ and a projective HT-drawing on S^2.

2.1 Hanani–Tutte drawings

The purpose of this section is to set up some of the conventions specific to the present work and introduce several tools used to manipulate with HT-drawings on the projective plane.

Let G be a graph and D be an HT-drawing of G on $\mathbb{R}P^2$. Let γ be the crosscap in the chosen representation of $\mathbb{R}P^2$. We note that the choice of γ is not unique. Later in this section we show that the choice of γ does not matter.

We consider a map $\lambda\colon E(G) \to \mathbb{Z}_2$. For an edge e, we let $\lambda(e)$ be the number of crossings of e and the crosscap γ modulo 2. We emphasize that λ depends on the choice of the crosscap. Afterwards, it will be useful to alter λ via so-called vertex-crosscap switches, which we will introduce later in the section.

Given a cycle Z in G, we can distinguish whether Z is drawn as a homologically non-trivial cycle by checking the value of $\lambda(Z) := \sum_{e\in E(Z)} \lambda(e) \in \mathbb{Z}_2$. The cycle Z is homologically non-trivial if and only if $\lambda(Z) = 1$. In particular, $\lambda(Z)$ does not depend on the choice of the crosscap (see 'Homology of surfaces' in 'Background and notation').

Let D be an HT-drawing of a graph G on $\mathbb{R}P^2$. It is not hard to derive a drawing D' of the same graph on S^2 such that every pair (e, f) of *independent* edges satisfies $\mathrm{cr}(e, f) = \lambda(e)\lambda(f)$. Indeed, it is sufficient to 'undo' the crosscap; that is, we look at the crosscap on S^2 used to represent $\mathbb{R}P^2$, undo the gluing, fill in the interior of the resulting disk B and then let the edges that passed over the crosscap intersect in B. See the two leftmost pictures in Figure 2.4. This motivates the following definition.

Definition 2.7 (Projective HT-drawings on S^2.)**.** *Let D be a drawing of a graph G on S^2 and $\lambda\colon E(G) \to \mathbb{Z}_2$ be a function. Then the pair (D, λ) is a* projective HT-drawing *of G on S^2 if* $\mathrm{cr}(e, f) = \lambda(e)\lambda(f)$ *for any pair of independent edges e and f of G. If λ is sufficiently clear from a context, we say that D is a projective HT-drawing of G on S^2.*

Clearly, an HT-drawing of G on $\mathbb{R}^2$ yields a projective HT-drawing of G on S^2. It turns out that it works in the opposite direction as well.

Lemma 2.8. *Let (D, λ) be a projective HT-drawing of a graph G on S^2. Then there is an HT-drawing $D_\otimes$ of G on $\mathbb{R}P^2$ such that $\mathrm{cr}_{D_\otimes}(e, f) = \mathrm{cr}_D(e, f) + \lambda(e)\lambda(f)$ for any pair of distinct edges of G, possibly adjacent. In addition, for*

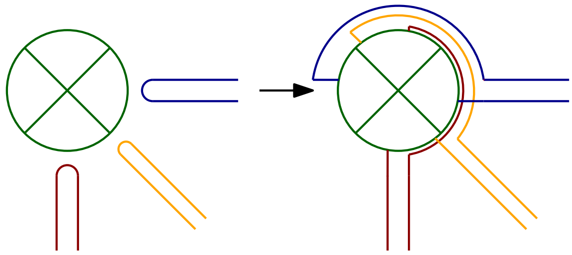

Figure 2.5: Redrawing the finger-moves around the crosscap.

any two edges e and f such that $\lambda(e) = \lambda(f) = 0$ and that $D(e)$ and $D(f)$ are disjoint, we have that $D_\otimes(e)$ and $D_\otimes(f)$ are disjoint as well.

Proof. It is sufficient to consider a small disk B which does not intersect $D(G)$, replace it with a crosscap and redraw the edges e with $\lambda(e) = 1$ appropriately as described below (follow the two pictures on the right in Figure 2.4). From each edge e with $\lambda(e) = 1$ we pull a thin 'finger'[11] towards the crosscap which intersects every other edge in pairs of intersection points. Then we redraw the edge in a close neighbourhood of the crosscap as indicated in Figure 2.5. After the redrawing, every edge e such that $\lambda(e) = 1$ passes over the crosscap once, while every edge e with $\lambda(e) = 0$ does not pass over the crosscap. This agrees with the original definition of λ for HT-drawings on $\mathbb{R}P^2$. In addition, we indeed obtain an HT-drawing on $\mathbb{R}P^2$ with $\mathrm{cr}_{D_\otimes}(e, f) = \mathrm{cr}_D(e, f) + \lambda(e)\lambda(f)$, because we have introduced one more crossing among pairs of edges e, f such that $\lambda(e) = \lambda(f) = 1$ in the last step. □

In summary, Lemma 2.8 together with the previous discussion provide us with two viewpoints on the HT-drawings.

Corollary 2.9. *A graph G admits a projective HT-drawing on S^2 with respect to a function $\lambda\colon E(G) \to \mathbb{Z}_2$ if and only if it admits an HT-drawing on $\mathbb{R}P^2$.*

The main benefit of the projective HT-drawings on S^2 lies in the fact that one can ignore the actual geometric position of the crosscap and work on the sphere instead, which is simpler. This is especially helpful when one needs to merge two drawings. On the other hand, it turns out that in the present approach it will be easier to perform certain parity counts in the language of HT-drawings on $\mathbb{R}P^2$.

In order to distinguish the usual HT-drawings on S^2 from the projective HT-drawings, we will sometimes refer to the former as to the *ordinary* HT-drawings on S^2.

Non-trivial walks. We now extend the notions of triviality/non-triviality from $\mathbb{R}P^2$ to the setting of projective HT-drawings on S^2.

Let (D, λ) be a projective HT-drawing of a graph G and ω be a walk in G. We define $\lambda(\omega) := \sum_{e \in E(\omega)} \lambda(e)$ where $E(\omega)$ is the multi-set of edges appearing in ω. Equivalently, it is sufficient to consider only the edges appearing an odd number of times in ω, since $2\lambda(e) = 0$ for any edge e. We say that ω is *trivial* if $\lambda(\omega) = 0$ and *non-trivial* otherwise.

[11] In the same way as we do it when performing a vertex-edge switch.

We often use this terminology in special cases when ω is an edge, a path or a cycle. In particular, a cycle Z is trivial if and only if it is drawn as a homologically trivial cycle in the corresponding drawing $D_\otimes$ of G on $\mathbb{R}P^2$ from Lemma 2.8.

Given two homologically non-trivial cycles on $\mathbb{R}P^2$, they must cross an odd number of times by Lemma 2.6. Since we have not proved Lemma 2.6 and only referred to literature, we prove here a weaker version of this statement in the setting of projective HT-drawings, which we will need soon.

Lemma 2.10. *Let (D, λ) be a projective HT-drawing of a graph G on S^2. Then G does not contain two vertex-disjoint non-trivial cycles.*

Proof. For contradiction, let Z_1 and Z_2 be two vertex-disjoint non-trivial cycles in G. This means that both Z_1 and Z_2 contain an odd number of non-trivial edges. Therefore, there is an odd number of pairs (e_1, e_2) of non-trivial edges such that $e_1 \in Z_1$ and $e_2 \in Z_2$. According to Definition 2.7, Z_1 and Z_2 must have an odd number of crossings in D. But this is impossible for two cycles in the plane that cross at every intersection. $\square$

Vertex-crosscap switches. Let (D, λ) be a projective HT-drawing of G on S^2. It is very useful to alter λ at the cost of redrawing G. Given a vertex v, we perform the vertex-edge switches (v, e) for all edges e not incident to v such that $\lambda(e) = 1$ obtaining a drawing D'. We also introduce a new function $\lambda' \colon E(G) \to \mathbb{Z}_2$ derived from λ by switching the value of λ on all edges of G incident to v. In this case, we say that D' (and λ') is obtained by the *vertex-crosscap switch* over v[12]. The result is again a projective HT-drawing:

Lemma 2.11. *Let (D, λ) be a projective HT-drawing of G on S^2. Let D' and λ' be obtained from D and λ by a vertex-crosscap switch. Then (D', λ') is a projective HT-drawing of G on S^2.*

Proof. We need to check that $\mathrm{cr}_{D'}(e, f) = \lambda'(e)\lambda'(f)$ for any pair of independent edges e and f. Let v be the vertex inducing the switch. If neither e nor f is incident to v, then

$$\mathrm{cr}_{D'}(e, f) = \mathrm{cr}_D(e, f) = \lambda(e)\lambda(f) = \lambda'(e)\lambda'(f).$$

Now consider the case that one of the edges, say e, is incident to v. Note that $\lambda(e) = 1 - \lambda'(e)$ and $\lambda(f) = \lambda'(f)$ in this case. If $\lambda(f) = 0$, then

$$\mathrm{cr}_{D'}(e, f) = \mathrm{cr}_D(e, f) = \lambda(e)\lambda(f) = 0 = \lambda'(e)\lambda'(f).$$

Finally, if $\lambda(f) = 1$, then

$$\mathrm{cr}_{D'}(e, f) = 1 - \mathrm{cr}_D(e, f) = 1 - \lambda(e)\lambda(f) = \lambda(f) - \lambda(e)\lambda(f) = \lambda'(e)\lambda'(f).$$

$\square$

We also remark that a vertex-crosscap switch preserves the triviality/non-triviality of cycles. To see this, consider a cycle Z. If Z avoids v, then $\lambda(Z) = \lambda'(Z)$ since $\lambda(e) = \lambda(e')$ for every edge e of Z. If Z contains v, then $\lambda(Z) = \lambda'(Z)$ as well since $\lambda(e) \neq \lambda'(e)$ for exactly two edges of Z.

[12]In the case of drawings on $\mathbb{R}P^2$, a vertex-crosscap switch corresponds to passing the crosscap over v, which motivated the name.

Planarization. As before, let (D, λ) be a projective HT-drawing of G on S^2. Now let us consider a subgraph P of G such that every cycle in P is trivial. Then P essentially behaves as a planar subgraph of G, which we make more precise in the following lemma.

Lemma 2.12. *Let (D, λ) be a projective HT-drawing of G on S^2 and let P be a subgraph of G such that every cycle in P is trivial. Then there is a set $U \subseteq V(P)$ with the following property. Let (D_U, λ_U) be obtained from (D, λ) by the vertex-crosscap switches over all vertices of U (in any order). Then (D_U, λ_U) is a projective HT-drawing of G on S^2 and $\lambda_U(e) = 0$ for every $e \in E(P)$.*

Proof. The drawing (D_U, λ_U) is a projective HT-drawing by Lemma 2.11. Let F be a spanning forest of P, the union of spanning trees of each connected component of P rooted arbitrarily. We first make $\lambda(e) = 0$ for each edge e of F as follows: we do a breadth-first search[13] on each tree in F starting in its root; whenever an edge $e \in F$ with $\lambda(e) = 1$ is encountered, we perform a vertex-crosscap switch on the endpoint of e farther from the root of the tree. Let λ_U be the resulting map, which is zero on the edges of F. Each edge e in $E(P) \setminus E(F)$ belongs to a cycle Z such that $Z - e \subseteq F$. Since $\lambda_U(Z) = \lambda(Z) = 0$, we have $\lambda_U(e) = 0$ as well. $\square$

[13]See, e.g., the book by Aho, Hopcroft, and Ullman [1].

2.2 Separation theorem

In this section, we state the separation theorem announced in 'Introduction' together with auxiliary tools needed for its proof. At the end, we also prove the separation theorem assuming the validity of the auxiliary tools, which will be proven in the remainder of the chapter.

As it was explained in 'Introduction', a simple cycle Z such that every edge of Z is even (in a drawing in the plane) splits the graph into the outside and the inside. We first introduce a notation related to this splitting.

Definition 2.13 (The inside and the outside graph)**.** *Let G be a graph and D be a drawing of G on S^2. Let us assume that Z is a cycle of G such that every edge of Z is even and it is drawn as a simple cycle in D. Let S^+ and S^- be the two components of $S^2 \setminus D(Z)$.*

We call a vertex $v \in V(G) \setminus V(Z)$ an inside *vertex if it belongs to S^+ and an* outside *vertex otherwise.*

Given an edge $e = uv \in E(G) \setminus E(Z)$, we say that e is an inside *edge if either u is an inside vertex, or if $u \in V(Z)$ and $D(e)$ points locally to S^+ next to $D(u)$. Analogously we define an* outside *edge.*[14]

We let V^+ and E^+ be the sets of the inside vertices and the inside edges, respectively. Analogously, we define V^- and E^-. We also define the graphs $G^{+0} := (V^+ \cup V(Z), E^+ \cup E(Z))$ and $G^{-0} := (V^- \cup V(Z), E^- \cup E(Z))$.

Now, we may formulate the main technical tool—the separation theorem for projective HT-drawings.

Theorem 2.14. *Let (D, λ) be a projective HT-drawing of a 2-connected graph G on S^2 and Z be a cycle of G that is simple in D and such that every edge of Z is even. Moreover, we assume that every edge e of Z is trivial, that is, $\lambda(e) = 0$. Then there is a projective HT-drawing (D', λ') of G on S^2 satisfying the following properties.*

- *The drawings D and D' coincide on Z.*
- *The cycle Z is completely free of crossings and all of its edges are trivial in D'.*
- *$D'(G^{+0})$ is contained in $S^+ \cup D'(Z)$.*
- *$D'(G^{-0})$ is contained in $S^- \cup D'(Z)$.*
- *All edges of G^{+0} or all edges of G^{-0} are trivial (according to λ'); that is, at least one of the drawings $D'(G^{+0})$ or $D'(G^{-0})$ is an ordinary HT-drawing on S^2.*

The assumption that G is 2-connected is not essential for the proof of Theorem 2.14, but it will slightly simplify some of the steps. For our application, it will be sufficient to prove the 2-connected case.

In the remainder of this section, we describe the main ingredients of the proof of Theorem 2.14 and we also derive the theorem using the ingredients. We will

[14] It turns out that every edge $e \in E(G) \setminus E(Z)$ is either an outside edge or an inside edge, since every edge of Z is even.

often encounter the setting when G, (D, λ) and Z satisfy the assumptions of Theorem 2.14:

Definition 2.15 (Separation assumptions)**.** *We say that G, (D, λ) and Z satisfy the* separation assumptions *if the following conditions are met at once:*

- *G is a 2-connected graph.*
- *(D, λ) is a projective HT-drawing of G.*
- *Z is a cycle in G drawn as a simple cycle in D.*
- *every edge of Z is even in D and trivial.*

Arrow graph. From now on, let us fix G, (D, λ) and Z satisfying the separation assumptions. Before we move on, let us recall a definition of a *bridge* from graph theory (see, e.g., Mohar and Thomassen [20, p. 7], where the bridges considered here are called 'Z-bridges'.)

Definition 2.16 (Bridges)**.** *A* bridge *B of G (with respect to Z) is a subgraph of G that is either an edge not in Z, but with both endpoints in Z (and its endpoints also belong to B), or a connected component of $G - V(Z)$ together with all edges (and their endpoints in Z) with one endpoint in that component and the other endpoint in Z.*

The distinction between the outside and the inside fixed by G, (D, λ) and Z satisfying the separation assumptions allows us to distinguish between the bridges inside and outside:

Definition 2.17 (Inside/outside bridges and proper walks)**.** *We say that B is an* inside bridge *if it is a subgraph of G^{+0}, and similarly, an* outside bridge *if it is a subgraph of G^{-0}*

A walk ω in G is a proper walk *if no vertex in ω belongs to $V(Z)$, except possibly its endpoints, and no edge of ω belongs to $E(Z)$. In particular, each proper walk belongs to a single bridge.*

As a consequence of the definitions, every bridge is either an inside bridge, or an outside bridge. Moreover, since we assume that G is 2-connected, every inside/outside bridge contains at least two vertices of Z. The bridges induce partitions of $E(G) \setminus E(Z)$ and of $V(G) \setminus V(Z)$. See Figure 2.6.

We want to record which pairs of vertices on $V(Z)$ are connected with a non-trivial and proper walk inside or outside[15]. In order to do this, we create two new multi-graphs A^+ and A^-, possibly with loops, but without multiple edges. In order to distinguish these graphs from G, we draw their edges with double arrows and we call them an *inside arrow graph* and an *outside arrow graph*, respectively. The edges of these graphs are called the *inside/outside arrows*. We set $V(A^+) = V(A^-) = V(Z)$.

Now we describe the *arrows*, that is, $E(A^+)$ and $E(A^-)$. Let u and v be two vertices of $V(Z)$, not necessarily distinct. By W^+_{uv} we denote the set of all proper, non-trivial walks in G^{+0} with endpoints u and v. We have an *inside arrow*

[15]We recall that non-trivial walks are defined in Section 2.1 below Corollary 2.9.

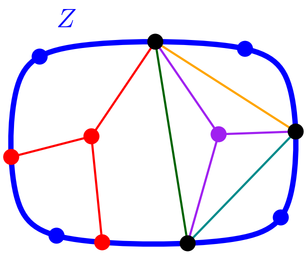

Figure 2.6: An example of a graph with five inside bridges—marked by different colours. The vertices that belong to several inside bridges are in black.

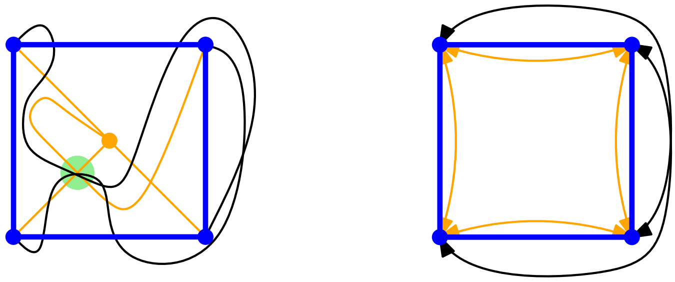

Figure 2.7: The inside and the outside arrows (right) corresponding to the projective HT-drawing of K_5 (left) derived from its drawing in Figure 2.2, left.

connecting u and v in $E(A^+)$ if and only if W^+_{uv} is non-empty. To distinguish the edges of G from the arrows, we denote an arrow by $\overline{uv} = \overline{vu}$. An arrow which is a loop at a vertex v is denoted by $\overline{vv}$. This often allows one to work with arrows $\overline{uv}$ without distinguishing between the cases $u = v$ or $u \neq v$. Analogously, we define the set W^-_{uv} and the *outside arrows*. See Figure 2.7 for the arrow graph(s) of the projective HT-drawing of K_5 corresponding to its drawing on $\mathbb{R}P^2$ depicted in Figure 2.2, left.

It follows from the definition of the inside bridges that any walk $\omega \in W^+_{uv}$ stays in one inside bridge. Given an inside bridge B, we let $W^+_{uv,B}$ be the set of all walks $\omega \in W^+_{uv}$ which belong to B. In particular, W^+_{uv} decomposes into the disjoint union of the sets $W^+_{uv,B_1}, \ldots, W^+_{uv,B_k}$ where $B_1, \ldots, B_k$ are all inside bridges. Given an inside arrow $\overline{uv}$ and an inside bridge B, we say that B *induces* $\overline{uv}$ if $W^+_{uv,B}$ is non-empty. (Note that an arrow can be induced by more than one bridge.) An inside bridge B is *non-trivial* if it induces at least one arrow. Given two inside arrows $\overline{uv}$ and $\overline{xy}$, we say that $\overline{uv}$ and $\overline{xy}$ *are induced by different bridges* if there are two different inside bridges B and B' such that B induces $\overline{uv}$ and B' induces $\overline{xy}$. As usual, we define analogous notions for the outside as well. Note that it may happen that there is an inside bridge inducing both $\overline{uv}$ and $\overline{xy}$ even if $\overline{uv}$ and $\overline{xy}$ are induced by different bridges.

Possible configurations of arrows. We plan to utilise the arrow graph in the following way. On the one hand, we will show that certain configurations of arrows are not possible; see Figure 2.8. On the other hand, we will show that, since the arrow graph does not contain any of the forbidden configurations, it has to contain one of the configurations in Figure 2.9 inside or outside. These

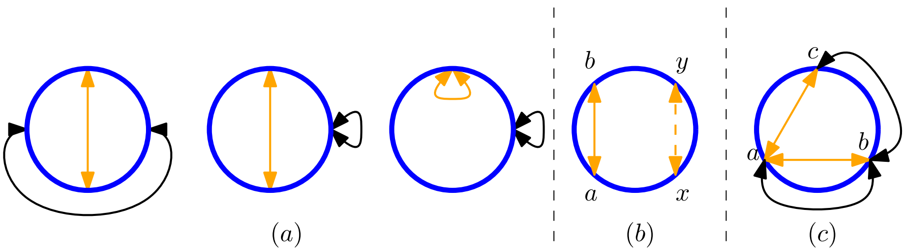

Figure 2.8: Forbidden configurations of arrows. The cyclic order along the cycle Z (in blue) in (a) may be arbitrary, whereas it is important in (b) that the arrows there do not interleave. Different dashing of lines in (b) correspond to arrows induced by different inside bridges. The arrows of the same colour in (c) are induced by the same bridge.

configurations are precisely defined in Definition 2.21 below.

We will also show that the configurations in Figure 2.9 are *redrawable*, that is, they may be redrawn without using the crosscap. The precise statement for redrawings is given by Proposition 2.23 below.

Specifically, we prove three lemmas forbidding the configurations of arrows depicted in Figure 2.8. We emphasize that in all these lemmas we assume that the notions used there correspond to a fixed G, (D, λ) and Z satisfying the separation assumptions.

Lemma 2.18. *Every inside arrow shares a vertex with every outside arrow.*

Every inside arrow shares a vertex with every outside arrow.

Lemma 2.19. *Let $\overline{ab}$ and $\overline{xy}$ be two arrows induced by different inside bridges of G^{+0}. If the two arrows do not share an endpoint, their endpoints have to interleave along Z.*

Lemma 2.20. *There are no three distinct vertices a, b, c on Z, an inside bridge B^+ and an outside bridge B^- such that B^+ induces the arrows $\overline{ab}$ and $\overline{ac}$ (and no other arrows) and B^- induces the arrows $\overline{ab}$ and $\overline{bc}$ (and no other arrows).*

We prove the three lemmas above in Section 2.5. By symmetry, Lemmas 2.19, 2.20 and 2.18 are also valid if we swap the inside and the outside.

We are ready to describe the redrawable configurations.

Definition 2.21 (Redrawable configurations of arrows)**.** *We say that G forms*

(a) *an* inside fan *if there is a vertex common to all inside arrows. The arrows may come from various inside bridges.*

(b) *an* inside square *if it contains four vertices a, b, c and d ordered in this cyclic order along Z and the inside arrows are precisely $\overline{ab}$, $\overline{bc}$, $\overline{cd}$ and $\overline{ad}$. In addition, we require that the inside graph G^{+0} has only one non-trivial inside bridge.*

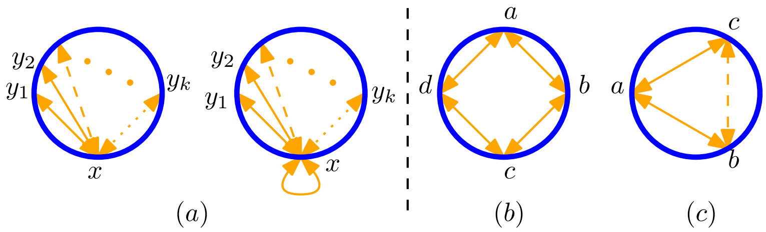

Figure 2.9: Schematic drawings of the redrawable configurations of arrows from Definition 2.21. Different dashing of lines correspond to different inside bridges. The loop in the right drawing (a) is an inside loop (drawn outside due to lack of space). The drawing (c) is only one instance of an inside split triangle.

(c) *an* inside split triangle *if there exist three vertices a, b, and c such that the inside arrows of G are $\overline{ab}$, $\overline{ac}$ and $\overline{bc}$. In addition, we require that every non-trivial inside bridge induces either the two arrows $\overline{ab}$ and $\overline{ac}$, or just a single arrow.*

See Figure 2.9. We use analogous definitions for an outside fan, outside square *and* outside split triangle.

We note that the notions in Definition 2.21 depend on G, (D, λ) and Z satisfying the separation assumptions.

A relatively direct case analysis using Lemmas 2.18, 2.19 and 2.20 reveals the following fact.

Proposition 2.22. *Let (D, λ) be a projective HT-drawing on S^2 of a graph G and let Z be a cycle in G satisfying the separation assumptions. Then G forms an (inside or outside) fan, square, or split triangle.*

On the other hand, any configuration from Definition 2.21 can be redrawn without using the crosscap:

Proposition 2.23. *Let (D, λ) be a projective HT-drawing of G^{+0} on S^2 and Z be a cycle satisfying the separation assumptions. Moreover, let us assume that $D(G^{+0}) \cap S^- = \emptyset$ (that is, G^{+0} is fully drawn on $S^+ \cup D(Z)$). Let us also assume that G^{+0} forms an inside fan, an inside square or an inside split triangle. Then there is an ordinary HT-drawing D' of G^{+0} on S^2 such that D coincides with D' on Z and $D'(G^{+0}) \cap S^- = \emptyset$.*

Proposition 2.22 is proven in Section 2.4 assuming there the validity of Lemmas 2.18, 2.19 and 2.20. Proposition 2.23 is proven in Section 2.6.

Now we are missing only one tool to finish the proof of Theorem 2.14—the 'redrawing procedure' of Pelsmajer et al. [22]. More precisely, we need the following variant of Theorem 2.5.

Theorem 2.24. *Let D be a drawing of a graph G on the sphere S^2. Let Z be a cycle in G such that every edge of Z is even and Z is drawn as a simple cycle. Then there is a drawing D'' of G such that*

- *D'' coincides with D on Z,*

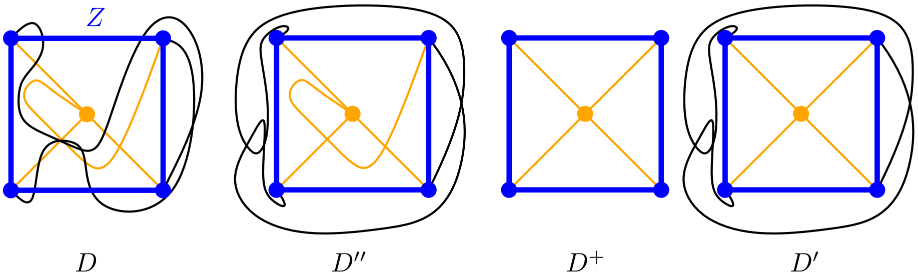

Figure 2.10: Redrawing a projective HT-drawing of K_5 analogously to the drawing in Figure 2.2.

- *$D''(G^{+0})$ belongs to $S^+ \cup D(Z)$ and $D''(G^{-0})$ belongs to $S^- \cup D(Z)$*
- *and whenever (e, f) is a pair of edges such that both e and f are inside edges or both e and f are outside edges, then $\mathrm{cr}_{D''}(e, f) = \mathrm{cr}_D(e, f)$.*

It is easy to check that the proof of Theorem 2.5 by Pelsmajer et al. [22] proves Theorem 2.24 as well. Additionally, we note that an alternative proof of Theorem 2.5 by Fulek et al. [10, Lem. 3] can also be extended to yield Theorem 2.24. Nevertheless, for completeness, we provide its proof in Section 2.7.

Finally, we prove the separation theorem (Theorem 2.14) assuming the validity of the aforementioned auxiliary results.

Proof of Theorem 2.14. Let G be the graph, (D, λ) be the drawing and Z be the cycle from the statement.

We use Theorem 2.24 with G and D to obtain a drawing D'' keeping in mind that all edges of Z are even. This is depicted in Figure 2.10, left; the following steps of the proof are illustrated there as well. We get that Z is drawn on D'' as a simple cycle free of crossings. We also get that $D''(G^{+0})$ is contained in $S^+ \cup D''(Z)$ and $D''(G^{-0})$ is contained in $S^- \cup D''(Z)$. However, there may be no λ'' for which (D'', λ'') would be a projective HT-drawing; we still may need to modify it to obtain such a drawing.

By Proposition 2.22 applied to (D, λ), G forms one of the redrawable configurations on one of the sides; that is, an inside/outside fan, square or split triangle. Without loss of generality, it appears inside. It means that D'' restricted to G^{+0} satisfies the assumptions of Proposition 2.23. Therefore, there is an ordinary HT-drawing D^+ of G^{+0} satisfying the conclusions of Proposition 2.23. Finally, we let D' be the drawing of G on S^2 which coincides with D^+ on G^{+0} and with D'' on G^{-0}. Both D'' and D^+ coincide with D on Z; therefore, D' is well defined. We set λ' so that $\lambda'(e) := \lambda(e)$ for an edge $e \in E^-$ and $\lambda'(e) := 0$ for any other edge.

Now we verify that (D', λ') is the required projective HT-drawing. Let e and f be independent edges. If both e and f are inside edges, then $\mathrm{cr}_{D'}(e, f) = \mathrm{cr}_{D^+}(e, f) = 0 = \lambda'(e)\lambda'(f)$, since D^+ is an ordinary HT-drawing. If both e and f are outside edges, then $\mathrm{cr}_{D'}(e, f) = \mathrm{cr}_{D''}(e, f) = \mathrm{cr}_D(e, f) = \lambda(e)\lambda(f) = \lambda'(e)\lambda'(f)$. Finally, if one of this edges is an inside edge and the other is an outside edge, then $\mathrm{cr}_{D'}(e, f) = 0 = \lambda'(e)\lambda'(f)$, because $D'(e)$ and $D'(f)$ are separated by $D'(Z)$. □

2.3 Proof of the strong Hanani–Tutte theorem on $\mathbb{R}P^2$

In this section we prove Theorem 2.4 assuming validity of Theorem 2.14 as well as some of the auxiliary results presented in the previous section, which will be proven only in the later sections.

Given a graph G that admits an HT-drawing on the projective plane, we need to show that G is actually projective-planar. By Corollary 2.9, we may assume that G admits a projective HT-drawing (D, λ) on S^2. We aim to use Theorem 2.14. For this, we need that G is 2-connected and contains a suitable trivial cycle Z that may be redrawn so that it satisfies the assumptions of Theorem 2.14. Therefore, we start with auxiliary claims that will bring us to this setting. Many of them are similar to auxiliary steps of Pelsmajer et al. [22]; sometimes they are almost identical, only adapted to a new setting.

Before we state the next lemma, we recall the well-known fact that any graph admits a unique decomposition into *blocks of 2-connectivity*; see, e.g., Diestel [7, Ch. 3]. Each block in the decomposition is either a vertex, an edge or a 2-connected graph with at least three vertices. Here, we also allow the case that G is disconnected. A block is a single vertex if and only if the vertex is isolated. The intersection of two blocks is either empty, or it contains a single vertex, which is a cut in the graph. The blocks of the decomposition cover all vertices and edges. A vertex may occur in several blocks, whereas each edge belongs to a unique block.

Lemma 2.25. *If G admits a projective HT-drawing on S^2, then at most one block of 2-connectivity in G is non-planar. Moreover, if all blocks are planar, G is planar as well.*

We note that Schaefer and Štefankovič [28] proved that a minimal counterexample to the strong Hanani–Tutte theorem on any surface is (vertex) 2-connected. However, for the projective plane the same result can be obtained by much simpler means; therefore, we include its proof here.

Proof. First, for contradiction, let us assume that G contains two distinct non-planar blocks B_1 and B_2. If B_1 and B_2 are disjoint, then Lemma 2.10 implies that at least one of these blocks, say B_2, does not contain any non-trivial cycle. However, it means that B_2 admits an ordinary HT-drawing on S^2 by Lemma 2.12. Therefore, B_2 is planar by the strong Hanani–Tutte theorem in the plane (Theorem 2.1). This contradicts our original assumption. It remains to consider the case when B_1 and B_2 share a vertex v (it must be a cut vertex). Let us set $H := B_1 \cup B_2$. Let P be a spanning tree of H with just two edges e_1, e_2 incident to v and such that $e_1 \in B_1$ and $e_2 \in B_2$. Note that such a tree always exists, because B_1 and B_2 are connected after removing v. By Lemma 2.12 we may assume that all the edges of P are trivial, after a possible alteration of λ.

Since P is a maximal subgraph of H not containing a cycle, any non-trivial edge e from $E(H) \setminus E(P)$ creates a non-trivial cycle in the corresponding block. If e is not incident to v, then the cycle avoids v by the choice of P. This means that, by Lemma 2.10, at least one of the blocks, say B_2, satisfies that all its non-trivial edges are incident with v. It follows that B_2 is a planar graph, because D is an

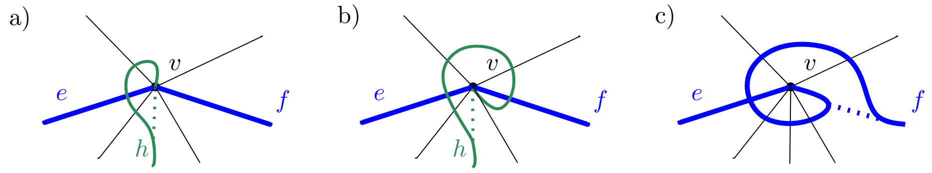

Figure 2.11: Local changes to make all edges of Z even. The original drawing of the edge near v is dotted.

HT-drawing of B_2 on S^2 (there is no pair of non-trivial independent edges in G). This is again a contradiction.

The last item in the statement of this lemma is a well-known property of planar graphs—a disjoint union of two planar graphs is a planar graph, and moreover, if a graph G contains a cut vertex v and all the components after cutting (and reattaching v) are planar, then G is planar as well. $\square$

Observation 2.26. *Let (D, λ) be a drawing of a 2-connected graph G. If D does not contain any trivial cycle, then G is planar.*

Proof. As G is 2-connected, it is either a cycle, or it contains three disjoint paths sharing their endpoints (see, e.g., Diestel [7, Prop. 3.1.2]). A cycle is a planar graph as we need. In the latter case, two of the paths are both trivial or both non-trivial. Together, they induce a trivial cycle, therefore this case cannot occur. $\square$

Lemma 2.27. *Let (D, λ) be a projective HT-drawing on S^2 of a graph G and let Z be a cycle in G. Then G can be redrawn only by local changes next to the vertices of Z to a projective HT-drawing D' on S^2 in such a way that λ remains unchanged and $cr_{D'}(e, f) = \lambda(e)\lambda(f)$ for any pair $(e, f) \in E(Z) \times E(G)$ of distinct (not necessarily independent) edges. In particular, if $\lambda(e) = 0$ for every edge e of Z, then every edge of Z becomes even in D'.*

Proof. Since we have a projective HT-drawing, $cr_D(e, f) = \lambda(e)\lambda(f)$ for every pair of independent edges. To prove the claim we need to show that local changes allow us to change the parity of $cr_D(e, f)$ whenever e is an edge of Z and e and f share a vertex.

This can be done in two steps. First, we use local move c) from Figure 2.11 to obtain the desired parity of $cr_D(e, f)$ for all pairs of consecutive edges (e, f) on Z. This move may change the parity of crossings between edges on Z and dependent edges not on Z.

Next, we use local moves a) and b) from Figure 2.11 to obtain the desired parity of crossings between edges on Z and dependent edges not on Z. If v is the vertex common to h, e and f, where e and f are edges on Z, move a) is used to change the parity of $cr_D(e, h)$ and its symmetric version to change the parity of $cr_D(f, h)$. Move b) is used to change the parity for both $cr_D(e, h)$ and $cr_D(f, h)$. Since these moves do not change the parity of $cr_D(e, h')$ or $cr_D(f, h')$ for any other edge h', the claim follows. $\square$

Once we know that the edges of a cycle can be made even, we also need to make such a cycle simple.

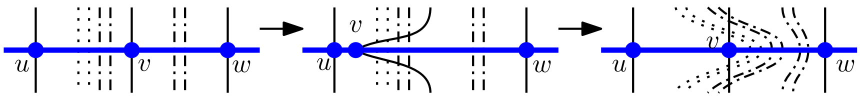

Figure 2.12: Almost contracting an edge.

Lemma 2.28. *Let (D, λ) be a projective HT-drawing on S^2 of a graph G and let Z be a cycle in G such that each of its edges is even. Then G can be redrawn so that Z becomes a simple cycle, its edges remain even and the resulting drawing is still a projective HT-drawing (with λ unchanged).*

Proof. First, we want to get a drawing such that there is only one edge of Z which may be intersected by other edges. Let us consider three consecutive vertices u, v and w on Z, with $v \notin \{u, w\}$. We *almost-contract* uv in the following way: we move the vertex v and the edges incident to v towards u until we remove all intersections between uv and other edges. Note that the image of the cycle Z is not changed; we only slide v towards u along the drawing of Z. This way, uv becomes free of crossings and its former crossings appear on vw instead. See the two leftmost pictures in Figure 2.12 (the right picture will be used in the proof of Theorem 2.24).

Since uv as well as vw were even edges in the initial drawing, vw remains even after the redrawing. Similarly, the parity of the number of crossings between the edges incident to v and other edges is not affected. If uv and vw intersected, then this step introduces self-intersections of vw.

After performing such redrawing repeatedly, we get that there is only one edge of Z which may be intersected by other edges, as required. We remove possible self-crossings of this edge and the other edges incident with v, as described in 'Background and notation'. □

Apart from lemmas tailored to set up the separation assumptions (see Definition 2.15), we also need one more lemma that will be useful in the inductive proof of Theorem 2.4.

Lemma 2.29. *Let $G, (D, \lambda)$ and Z satisfy the separation assumptions (see Definition 2.15). Let B be an inside bridge such that every proper path in B with both endpoints on $V(B) \cap V(Z)$ is non-trivial. Then $|V(B) \cap V(Z)| = 2$ and B induces a single arrow and no loop.*

Proof. First, we show that there is no non-trivial cycle in B. For contradiction, there is a non-trivial cycle N in B. By the 2-connectivity of G there exist two vertex disjoint paths p_1 and p_2, possibly of length zero, that connect Z to N. We consider shortest such paths; thus, each of the paths shares only one vertex with Z and one vertex with N. Let y_1 and y_2 be the endpoints of p_1 and p_2 on N, respectively. Let p_3, p_4 be the two arcs of N between y_1 and y_2. We consider two paths q_1 and q_2: q_1 is obtained from the concatenation of p_1, p_3 and p_2, while q_2 is defined as the concatenation of p_1, p_4 and p_2. Since N is non-trivial, one of these paths has to be trivial, which provides the required contradiction.

Next, we observe that B does not induce any loop in the inside arrow graph A^+. For contradiction, it induces a loop at a vertex $x \in V(Z)$. This means that there is a closed, proper and non-trivial walk κ in B containing x. We choose

κ as the shortest such walk. We already know that κ cannot be a cycle, thus it contains a closed non-empty subwalk κ'. Again, we choose κ' as the shortest such subwalk. Therefore, κ' must be a cycle; by the previous part of this proof, it is trivial. However, it means that κ can be shortened by leaving out κ', which is the required contradiction.

Now we show that $|V(B) \cap V(Z)| = 2$. By the 2-connectedness of G, we have that $|V(B) \cap V(Z)| \geq 2$. Thus, for contradiction, let a, b, c be three distinct vertices of $V(B) \cap V(Z)$. Let v be one of the inner vertices of B; there must be such a vertex, since B cannot be a single edge in this case. By the definition of inside/outside bridges, there exist proper walks p_a, p_b and p_c connecting v to a, b and c, respectively. By the pigeonhole principle, two of the walks have the same value of λ; without loss of generality, let $\lambda(p_a) = \lambda(p_b)$. Hence, the proper walk obtained from the concatenation of p_a and p_b is trivial. Since B does not contain any non-trivial cycle, this walk can be shortened to a trivial proper path between a and b by an argument analogous to the one used in the previous paragraph; a contradiction.

Finally, we know that there are two vertices in $V(B) \cap V(Z)$; let us write x and y for them. Since any path connecting x and y is non-trivial, B induces the arrow $\overline{xy}$ in A^+. No other arrow in A^+ induced by B is possible, since there are no loops. $\square$

Proposition 2.30 below is the main tool for deriving Theorem 2.4 from Theorem 2.14. It is set up in a way suitable for an induction based on Theorem 2.14.

Proposition 2.30. *Let (D, λ) be a projective HT-drawing of a 2-connected graph G on S^2 and Z be a cycle in G that is completely free of crossings in D and such that each of its edges is trivial in D. Assume that (V^+, E^+) or (V^-, E^-) is empty (recall the notation from Definition 2.13). Then G can be embedded into $\mathbb{R}P^2$ so that Z bounds a face in the resulting embedding homeomorphic to a disk. If, in addition, D is an ordinary HT-drawing on S^2, then G can be embedded into S^2 so that Z bounds a face of the resulting embedding*[16].

We need to consider the case of ordinary HT-drawings in this proposition for a well-working induction.

Proof. The proof proceeds by induction on the number of edges of G. The base case is when G is a cycle.

Without loss of generality, we assume that (V^-, E^-) is empty, i.e., $G = G^{+0}$. If (V^+, E^+) is also empty, G consists only of Z and such a graph can easily be embedded into the sphere or the projective plane as required. Therefore, we assume that (V^+, E^+) is non-empty.

We find a path γ in $(V(G^{+0}), E(G^{+0}) \setminus E(Z))$ connecting two points x and y lying on Z. We may choose x, y so that $x \neq y$, since G is 2-connected. There are two cases to consider.

Case 1: There exists a trivial γ. First we solve the case that at least one such path γ is trivial. We show that all edges of γ can be made even and simple

[16]This face is again homeomorphic to a disk—there is, in fact, no other option on S^2.

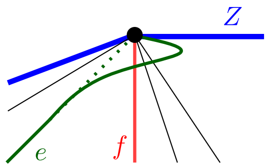

Figure 2.13: Local changes at u. The original drawing of the edge is dotted, Z is depicted in blue, f (as a part of γ) in red. The changed edge in green.

in the drawing, while keeping Z simple and free of crossings in the drawing and the whole drawing stays a projective HT-drawing of G^{+0}.

As the first step, we use Lemma 2.12 in order to ensure that $\lambda(e) = 0$ for every edge $e \in E(Z) \cup E(\gamma)$. Inspecting the proof of Lemma 2.12 we see that this can be achieved by performing vertex-crosscap switches only over the inner vertices of γ—it is sufficient to set up the root in the proof of Lemma 2.12 to be one of the endpoints of γ (the edges of Z are already trivial by assumption). We can perform the vertex-crosscap switches in S^+, which is determined by Z, without affecting Z.

Now we want to make the edges of γ even, but again without affecting Z. First, for every pair (e, f) of adjacent edges of γ which intersect oddly, we locally perform the move c) from Figure 2.11 similarly as in Lemma 2.27. Next, we consider an edge $e \notin E(\gamma)$ adjacent to a vertex $u \in V(\gamma) \setminus V(Z)$. For such an edge we can perform one of the moves a) or b) depicted in Figure 2.11, if needed, so that e then intersects evenly each of the two edges of γ incident with u. Finally, we consider an edge $e \notin E(\gamma) \cup E(Z)$ adjacent to $u \in \{x, y\}$, one of the endpoints of γ on Z. Let f be the edge of γ incident with u. If e and f intersect oddly, we perform the move from Figure 2.13. This is possible since Z is free of crossings. This way we achieve that every edge of γ becomes even.

As the last step in redrawing γ, we want to make γ simple (again without affecting Z). This can be done in the same way as in the proof of Lemma 2.28. We almost-contract all but on edge of γ so that there remains only one edge of γ that may intersect other edges. Then we remove possible self-intersections, as described in 'Background and notation'.

The rest of the argument is easier to explain if we switch inside and outside (this is easily doable by a homeomorphism of S^2) and treat drawings on S^2 as drawings in the plane.

We may assume that, after a homeomorphism, Z is drawn in the plane as a geometric circle with its inner region empty and with x and y antipodal[17]. The vertices x and y split Z into two paths; we denote by p_1 the 'upper' one and by p_2 the 'lower' one. We may also assume that γ is 'above' p_1 by adapting the initial choice of the correspondence between S^2 and the plane, if necessary.

Now we continuously deform the plane in such a way that Z becomes flatter and flatter until it coincides with the line segment connecting x to y, as depicted in Figure 2.14, part a). We may further require that no inner vertex of p_1 was identified with an inner vertex of p_2.

This way, we get a projective HT-drawing $(\bar{D}, \bar{\lambda})$ of a new graph $\bar{G}$: all the vertices of G remain present in $\bar{G}$, that is, $V(G) = V(\bar{G})$. Also the edges of G

[17]That is, x and y are the opposite points on the circle representing Z.

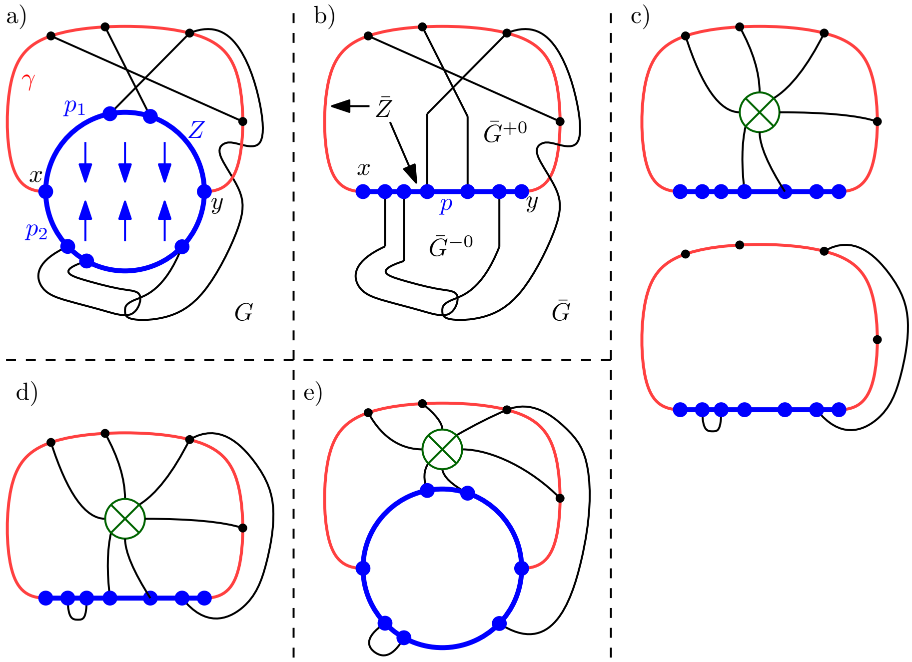

Figure 2.14: The deformation of the plane that changes G into $\bar{G}$, the redrawing of $\bar{G}$ and the resulting embeddings of $\bar{G}$ and G.

which are not on Z are present in $\bar{G}$. Only some of the edges of Z may disappear and they are replaced with edges forming a path p between x and y. Note that we have not introduced any parallel edges, because there is no edge in G connecting an inner vertex of p_1 with an inner vertex of p_2 (such an edge would have to cross γ oddly). It also turns out that $\bar{G}$ has one edge less than G. Regarding $\bar{\lambda}$, we have $\lambda(e) = \bar{\lambda}(e)$ if $e \in E(G) \setminus E(Z)$ and we have $\bar{\lambda}(e) = 0$ if $e \in E(p)$.

Now consider the cycle $\bar{Z}$ in $\bar{G}$ formed by γ and p. It is trivial and simple. In particular, it splits the plane into the inside and the outside as in Definition 2.13. For example, $\bar{G}^{+0}$ corresponds to the part of G in between γ and p_1 before the flattening; see Figure 2.14, parts a) and b).

Next, we apply Theorem 2.14 and get a drawing D' of $\bar{G}$. Looking at the two sides of $\bar{G}$ separately, we see a projective HT-drawing on one of the sides and an ordinary HT-drawing on S^2 on the other side. If, in addition, D was already an ordinary HT-drawing, we get an ordinary HT-drawing on both sides by Theorem 2.24. Without loss of generality, we assume that $\bar{G}^{+0}$ is the part that may be drawn as a projective HT-drawing.

Note also that since G was 2-connected, both parts of $\bar{G}$ are 2-connected as well. Subsequently, we examine each of these two parts separately and use the inductive hypothesis; we obtain an embedding of $\bar{G}^{+0}$ into $\mathbb{R}P^2$ such that $\bar{Z}$ bounds a face homeomorphic to a disk as well as an embedding of $\bar{G}^{-0}$ into S^2 such that $\bar{Z}$ bounds a face homeomorphic to a disk. If, in addition, D was already an ordinary HT-drawing, we get also the required embedding of $\bar{G}^{+0}$ into S^2. We merge the two embeddings along $\bar{Z}$ and obtain an embedding of $\bar{G}$ into $\mathbb{R}P^2$ (or S^2 if D was an ordinary HT-drawing). See Figure 2.14, parts c) and d).

Finally, we need to undo the identification of p_1 and p_2, which created p. Whenever we consider a vertex $v \in V(p)$ different from x and y, it is uniquely determined whether it comes from p_1 or p_2. In addition, if v comes from p_1, then every edge $e \in E(G) \setminus E(Z)$ incident with v must belong to $\bar{G}^{+0}$. Similarly, if v comes from p_1, then each edge $e \in E(G) \setminus E(Z)$ incident with v must belong to $\bar{G}^{-0}$. Therefore, it is possible to undo the identification and get the required embedding of G. See Figure 2.14, part e).

Case 2: All choices of γ are non-trivial. Now we deal with a situation in which all possible choices of γ are non-trivial. Let us consider the inside arrow graph A^+. Since all choices of γ are non-trivial, Lemma 2.29 shows that every inside bridge induces a single inside arrow. This will allow us to redraw inside bridges separately in a way described in the following claim.

Claim 2.30.1. *For every inside bridge B there exists a planar drawing of $Z \cup B$ in which Z is the outer face.*

Proof. Since we know that B induces only a single arrow, we get that $Z \cup B$ forms an inside fan according to Definition 2.21. It follows from Proposition 2.23 that $Z \cup B$ admits an ordinary HT-drawing such that Z is an outer cycle. However, the setting of ordinary HT-drawings is already fully resolved in Case 1. That is, we may already use Proposition 2.30 for this drawing and we get the required conclusion. □

Consider the graph A^{+0} obtained by adding the edges of Z to A^+, where A^+ is the inside arrow graph. Note that $V(A^+) = V(Z)$ according to the definition of the arrow graph.

Now our aim will be to find an embedding of A^{+0} into $\mathbb{R}P^2$ in which Z bounds a face. Once we achieve this, we can replace an embedding of each arrow by the embedding of inside bridges inducing this arrow via Claim 2.30.1 in a close neighbourhood of the arrow. If there are more inside bridges inducing the arrow, they are embedded in parallel.

Finally, we show that it is possible to embed A^{+0} in the required way. By Lemma 2.19, any two disjoint arrows interleave.

Let us consider two concentric closed disks E_1 and E_2 such that E_1 belongs to the interior of E_2. Let us draw Z on the boundary of E_1. Let a be the number of arrows of A^+ and let us consider $2a$ points on the boundary of E_1 placed to form the vertices of a regular $2a$-gon. These points will be marked by ordered pairs (x, y) where $\overline{xy}$ is an inside arrow. We mark the points in such a way that the cyclic order of the vertices from the first co-ordinate of the marks on the points respects the cyclic order of $V(Z)$ along Z. In particular, marks with the same first co-ordinate are consecutive. However, for a fixed x, the marks $(x, y_1), \dots (x, y_k)$ corresponding to all arrows emanating from x are ordered in the reversed order when compared with the order of $y_1, \dots, y_k$ along Z. See Figure 2.15.

We show that the points marked (x, y) and (y, x) have to be the opposite points on E_1 for every inside arrow $\overline{xy}$. For contradiction, let us assume that (x, y) and (y, x) are not opposite for some $\overline{xy}$. Then there is another arrow $\overline{uv}$ such that the marks (x, y) and (y, x) do not interleave with (u, v) and (v, u)—such an arrow exists, because the arrows induce a pairing of the vertices of E_1, and (x, y) and (y, x) do not split the vertices of E_1 equally. However, if $\overline{xy}$ and $\overline{uv}$ do not share an endpoint, we get a contradiction with the fact that disjoint arrows

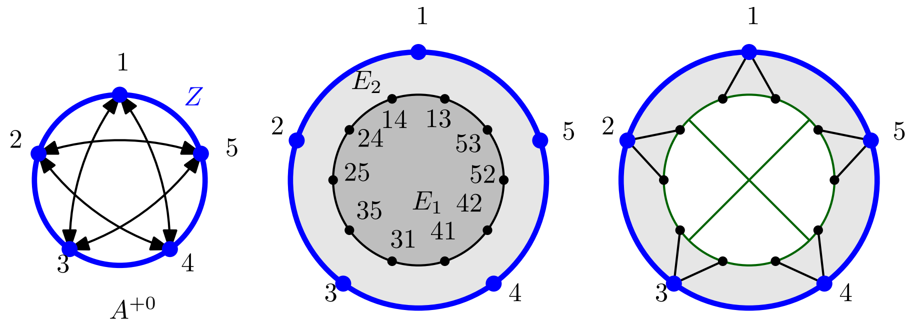

Figure 2.15: Redrawing the case in which every inside bridge induces a single arrow.

interleave; see Lemma 2.19. If $\overline{xy}$ and $\overline{uv}$ share an endpoint, we get a contradiction with the fact that we have reversed the order in the second co-ordinate of marks.

Now we get the required drawing, as depicted in Figure 2.15, in the following way. For an arrow $\overline{xy}$ we connect x with the point (x, y), and similarly, y with (y, x). We can add all such connections simultaneously for all arrows without introducing any crossings, since we have respected the cyclic order of $V(Z)$ along Z in the first co-ordinate of marks. We remove the interior of E_1 and we identify the pairs of opposite points on the boundary. This way we introduce a crosscap (see 'Background and notation'). Finally, we glue another disk along the boundary of E_2 to get a sphere with a crosscap; in other words, we get the required drawing on $\mathbb{R}P^2$. $\square$

Finally, we can prove Theorem 2.4.

Proof of Theorem 2.4. We prove the result by induction in the number of vertices of G. We can trivially assume that G has at least three vertices. We consider two cases depending on the connectedness of G.

If G has at least two blocks of 2-connectivity, G can be written as $G_1 \cup G_2$, where $G_1 \cap G_2$ is a minimal cut of G and, therefore, has at most one vertex. By Lemma 2.25, we may assume that G_1 is planar and G_2 non-planar. By induction, there exists an embedding D_2 of G_2 into $\mathbb{R}P^2$. Consequently, G_1 is planar, G_2 is embeddable into $\mathbb{R}P^2$ and $G_1 \cap G_2$ has at most one vertex. From these two embeddings we can easily derive an embedding of $G = G_1 \cup G_2$ in $\mathbb{R}P^2$.

We are left with the case that G is 2-connected. By Observation 2.26, we may assume that there is at least one trivial cycle Z in (D, λ). We can also make each of its edges trivial by Lemma 2.12 and even by Lemma 2.27. In addition, we make Z simple using Lemma 2.28. Hence G, Z and the current projective HT-drawing satisfy the separation assumptions (see Definition 2.15).

Now we use Z to redraw G as follows. First, we apply Theorem 2.14 to get a projective HT-drawing (D', λ') that separates G^{+0} and G^{-0}. We define $D^+ := D'(G^{+0})$ and $D^- := D'(G^{-0})$. Without loss of generality, D^- is an ordinary HT-drawing on S^2, while D^+ is a projective HT-drawing on S^2. Next, we apply Proposition 2.30 established above to D^+ and D^- separately. Thus, we get embeddings of G^{+0} and G^{-0}—one of them in S^2, the other one in $\mathbb{R}P^2$. In addition, Z bounds a face in both of them; hence, we can easily glue them to get an embedding of the whole graph G into $\mathbb{R}P^2$. $\square$

2.4 Labellings of inside/outside bridges and the proof of Proposition 2.22

In this section, given an inside (or outside) bridge B, we first describe what are possible combinations of arrows induced by B. Then we use the obtained findings to prove Proposition 2.22, assuming validity of Lemmas 2.18, 2.19 and 2.20, which will be proved in Section 2.5.

Labelling the vertices of the bridges. To describe possible combinations of arrows induced by a bridge B, we introduce certain *labellings* of $V(B) \cap V(Z)$.

Definition 2.31 (Labelling of $V(B) \cap V(Z)$)**.** *Let B be a bridge. A* valid *labelling $L = L_B$ for B is a mapping $L\colon V(B) \cap V(Z) \to \{\{0\}, \{1\}, \{0,1\}\}$ obtained in the following way.*

If $V(B) \setminus V(Z) \neq \emptyset$ we pick a reference vertex *$v_B \in V(B) \setminus V(Z)$ for L. Then we fix a* labelling parameter *$\alpha_B \in \mathbb{Z}_2$ for L. Finally, for any $u \in V(B) \cap V(Z)$ and for any proper walk ω with endpoints u and v_B, the vertex u receives the label $\alpha_B + \lambda(\omega) \in \mathbb{Z}_2$. Note that u may receive two labels after considering all such walks. On the other hand, each vertex of $V(B) \cap V(Z)$ obtains at least one label, which follows from the definition of bridges (Definition 2.16).*

If $V(B) \subseteq V(Z)$, then B comprises of only one edge $e = uv$ connecting two vertices of $V(Z)$. In this case, there are two valid labellings for B. We set $L(u) = \{\alpha_B\}$ and $L(v) = \{\lambda(e) + \alpha_B\}$ for a chosen labelling parameter $\alpha_B \in \mathbb{Z}_2$. We do not define a reference vertex in this case.

We often refer to a valid labelling only as to *labelling*. We use the word *valid* when we want to emphasize that we do not consider arbitrary assignment of zeroes and ones to $V(B) \cap V(Z)$.

If the bridge B is understood from the context, we may write just v instead of v_B for the reference vertex and α instead of α_B for the labelling parameter. Switching the choice of α in the definition we swap all labels. This means that there are at least two valid labellings for a given bridge, unless all the labels are $\{0,1\}$. On the other hand, as will be explained below, a different choice of the reference vertex either does not influence the resulting labelling, or has the same effect as swapping the value of the labelling parameter α.

To see this, consider a vertex $u \in V(B) \setminus V(Z)$ different from $v = v_B$. By Definition 2.16, there is a proper uv-walk γ in B, which has to be disjoint from Z. Now for any $x \in V(B) \cap V(Z)$ and every proper xv-walk ω_{xv} in B, the concatenation of the walks ω_{xv} and γ is a proper xu-walk in B of type $\lambda(\omega_{xv}) + \lambda(\gamma)$. Moreover, for any proper xu-walk ω_{xu} in B, the concatenation of the walks ω_{xu} and γ is a proper xv-walk in B of type $\lambda(\omega_{xu}) + \lambda(\gamma)$. As a result, choosing u as the reference vertex with $\alpha + \lambda(\gamma)$ as the labelling parameter leads to the same labelling as the choice of v as the reference vertex with the labelling parameter α. This idea can be used to establish the following simple, but useful observation.

Observation 2.32. *Let B be a bridge such that $V(B) \setminus V(Z) \neq \emptyset$. Moreover, let L be a valid labelling for B and v the reference vertex for L. Let $x, y \in V(B)$ and ω be a proper xy-walk in B. Then there is a proper xy-walk ω' in B containing the reference vertex v such that $\lambda(\omega) = \lambda(\omega')$.*

Proof. If ω contains inside/outside vertices, we choose one of them and denote it by u.

In the case that ω does not contain any such vertex, then $x \in V(Z)$ and $x = y$, since B cannot consist of just one edge. That is, ω is the walk of length zero with the single vertex x. In this case we choose $u = x$.

Now we find a proper uv-walk γ in B and use it as a detour. More precisely, ω' starts at x and follows ω to the first occurrence of u in ω. Then it goes to v and back along γ. Finally, it continues to y along ω. It is clear that $\lambda(\omega) = \lambda(\omega')$. By the choice of u, the walk ω' is also proper. □

The reason to introduce the labellings lies in the following property. Let L_B be a valid labelling of a bridge B. For any $u, w \in V(B) \cap V(Z)$, there is an arrow $\overline{uw}$ induced by B if and only if the vertices u and w were assigned different labels by L_B—this is proved in the next proposition. As usual, we state and prove the proposition only for inside bridges, but the same result holds for outside bridges analogously.

Proposition 2.33. *Let B be an inside bridge and L be a valid labelling for B. Let $x, y \in V(B) \cap V(Z)$ (possibly $x = y$). Then the inside arrow graph A^+ contains an arrow $\overline{xy}$ arising from B if and only if $L(x) \cup L(y) = \{0, 1\}$.*

Proof. It is straightforward to check the claim if B is just an edge e. Indeed, if $x \neq y$, then $e = xy$, and it defines the arrow $\overline{xy}$ arising from B if and only if $\lambda(e) = 1$, which in turn happens if and only if $L(x) \cup L(y) = \{0, 1\}$ according to Definition 2.31. If $x = y$, then $\overline{xx}$ is not induced by B and, at the same time, $|L(x) \cup L(x)| = 1$.

If $V(B) \setminus V(Z) \neq \emptyset$, let $v = v_B$ be the reference vertex for L. First, let us assume that $L(x) \cup L(y) = \{0, 1\}$. Let us consider a proper xv-walk ω_{xv} and a proper vy-walk ω_{vy} in B such that $\lambda(\omega_{xv}) \neq \lambda(\omega_{vy})$. Such walks exist by Definition 2.31, since $L(x) \cup L(y) = \{0, 1\}$. Then the concatenation of ω_{xv} and ω_{vy} yields a non-trivial walk, which belongs to $W^+_{xy,B}$; therefore, $\overline{xy}$ is induced by B.

On the other hand, let us assume that there is a non-trivial walk ω in $W^+_{xy,B}$ defining the arrow $\overline{xy}$. We know that ω is not just an edge, because it would mean that B consisted only of that edge. By Observation 2.32, we may assume that ω contains the reference vertex v. This vertex splits ω into two proper walks ω_1 and ω_2, each containing at least one edge. Since $\lambda(\omega) = 1$, we have $\lambda(\omega_1) \neq \lambda(\omega_2)$. Consequently, $L(x) \cup L(y) = \{0, 1\}$. □

The argument from the last two paragraphs of the proof above can also be used to establish the following lemma.

Lemma 2.34. *Let B be a bridge and L be a valid labelling for B. Moreover, let $x, y \in V(B) \cap V(Z)$ be two distinct vertices. If $|L(x)| = |L(y)| = 1$, then for any two proper xy-walks ω_1, ω_2 in B we have $\lambda(\omega_1) = \lambda(\omega_2)$.*

Proof. If B contains just the edge xy, the observation is trivially true. Therefore, we assume that there is the inside/outside reference vertex $v \in V(B)$ for L. By the assumption, every two proper xv-walks in B have the same λ-value. The same holds also for proper vy-walks in B. By Observation 2.32, we can assume that both ω_1 and ω_2 contain v. Then the lemma follows. □

Another useful consequence of Proposition 2.33 is the following description of inside arrows induced by an inside bridge which does not induce any loop. Similarly, it holds for the outside.

Lemma 2.35. *Let B be an inside bridge which does not induce any loop. Then the inside arrows induced by B form a complete bipartite graph. (One of the parts is empty if B does not induce any arrow.)*

Proof. Let us consider a valid labelling L for B. By Proposition 2.33, $|L(x)| = 1$ for any $x \in V(B) \cap V(Z)$, since B does not induce any loop. By Proposition 2.33 again, the inside arrows induced by B form a complete bipartite graph, in which one part corresponds to the vertices labelled 0 and the second part corresponds to the vertices labelled 1. □

We conclude this section a by a proof of Proposition 2.22, which asserts that, given $G, (D, \lambda)$ and Z satisfying the separation assumptions, the graph G forms an (inside or outside) fan, square, or split triangle in the drawing D.

Proof of Proposition 2.22. We need to distinguish a few cases.

First, we consider the case when we have two disjoint inside arrows, but at least one of them is a loop. In this case, it is easy to see that Lemma 2.18 implies that G forms the outside fan and we are done; see Figure 2.16, part a).

Second, let us consider the case that we have two disjoint inside arrows $\overline{ab}$ and $\overline{cd}$ which are not loops. Lemma 2.18 implies that the only possible outside arrows are $\overline{ac}$, $\overline{ad}$, $\overline{bc}$, $\overline{bd}$. (In particular, there are no loops outside.) If there are not two disjoint arrows outside, then G forms an outside fan and we are done. Therefore, we may assume that there are two disjoint arrows outside, without loss of generality, $\overline{ac}$ and $\overline{bd}$ (otherwise we swap a and b). By swapping outside and inside in the previous argument, we get that only further possible arrows inside are $\overline{ad}$ and $\overline{bc}$.

Now we distinguish a subcase when there is an inside bridge inducing the inside arrows $\overline{ab}$ and $\overline{cd}$. In this case, $\overline{ad}$ and $\overline{bc}$ have to be inside arrows as well by Lemma 2.35. By Lemma 2.18, we know that $\overline{ac}$ and $\overline{bd}$ are the only outside arrows; in particular, they are induced by different outside bridges by a variant of Lemma 2.35 for the outside. By Lemma 2.19 we see that they have to alternate. That is, up to renaming of the vertices, we get the right cyclic order for an inside square; see Figure 2.16, part b). In order to check that G indeed forms an inside square, it remains to verify that G has only one non-trivial inside bridge. The inside arrows are $\overline{ab}$, $\overline{bc}$, $\overline{cd}$ and $\overline{ad}$. If any of these arrows, for example $\overline{ab}$, is induced by two bridges, then we get a contradiction with Lemma 2.19, in this case on arrows $\overline{ab}$ and $\overline{cd}$.

By swapping inside and outside we solve the subcase when there is an outside bridge inducing the outside arrows $\overline{ac}$ and $\overline{bd}$; we get that G forms an outside square.

It remains to consider the subcase when $\overline{ab}$ and $\overline{cd}$ arise from different inside bridges and $\overline{ac}$ and $\overline{bd}$ arise from different outside bridges. However, Lemma 2.19 applied to the inside and then to the outside reveals that these two events cannot happen simultaneously.

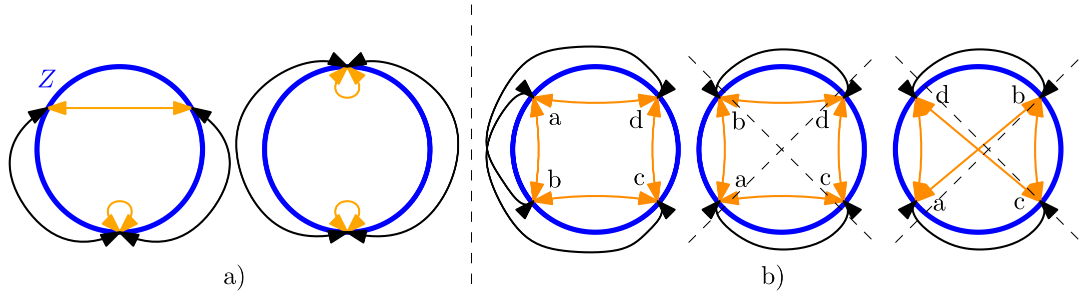

Figure 2.16: Schematic drawings of possible configurations of arrows. In a) only some of the outside arrows may be present. In b) all the arrows in the first figure must be present. The two crossed-out drawings are not possible.

Consequently, we have proved Proposition 2.22 in the case that there are two disjoint inside arrows. Analogously, we resolve the case when we have two disjoint arrows outside.

Finally, we consider the case when every pair of inside arrows shares a vertex and every pair of outside arrows shares a vertex. If there is a vertex v common to all the inside arrows, then we get an inside fan and we are done.

It remains to consider the last subcase when there is no vertex common to all inside arrows while every pair of inside arrows shares a vertex. This leaves the only option that there are three distinct vertices a, b and c on Z and all three inside arrows $\overline{ab}, \overline{ac}$ and $\overline{bc}$ are present; see Figure 2.9, part c). Then the only possible outside arrows are $\overline{ab}, \overline{ac}$ and $\overline{bc}$ as well due to Lemma 2.18. In addition, we may assume that all three outside arrows $\overline{ab}$, $\overline{ac}$ and $\overline{bc}$ are present; otherwise we have an outside fan and we are done.

In the present case, an inside bridge can induce at most two arrows by Lemma 2.35. Let us consider the three pairs of arrows $\{\overline{ab}, \overline{ac}\}$, $\{\overline{ab}, \overline{bc}\}$, and $\{\overline{ac}, \overline{bc}\}$. If at most one of these pairs is induced by an inside bridge, then G forms an inside split triangle and we are done. Analogously, we are done, if at most one of these pairs is induced by an outside bridge. Therefore, it remains to consider the case that at least two such pairs are induced by inside bridges and at least two such pairs are induced by outside bridges. However, this contradicts Lemma 2.20; see Figure 2.8, part c) as well. □

2.5 Forbidden configurations of arrows

In this section we show that certain combinations of arrows are not possible. That is, we prove Lemmas 2.18, 2.19 and 2.20. As before, we have a fixed graph G, its drawing (D, λ) on S^2 and a cycle Z in G. Again, we assume that $G, (D, \lambda)$ and Z satisfy the separation assumptions.

From the sphere to the projective plane. Although it is overall simpler to present the proof of Theorem 2.4 in the setting of projective HT-drawings on S^2, it is easier to prove Lemmas 2.18, 2.19 and 2.20 in the setting of HT-drawings on $\mathbb{R}P^2$. A small drawback is that we need to check that splitting of S^2 to the inside and outside part works for $\mathbb{R}P^2$ analogously.

Lemma 2.36. *Let (D, λ) be a projective HT-drawing of a graph G on S^2 and let Z be a cycle satisfying the separation assumptions. Let $D_\otimes$ be the HT-drawing of G on $\mathbb{R}P^2$ coming from the proof of Lemma 2.8. Then $D_\otimes(Z)$ is a simple cycle such that each of its edges is even, which separates $\mathbb{R}P^2$ into two parts, $(\mathbb{R}P^2)^+$ and $(\mathbb{R}P^2)^-$. In addition, every inside edge (with respect to D) which is incident to a vertex of Z points locally into $(\mathbb{R}P^2)^+$ in $D_\otimes$ as well as every outside edge (with respect to D) which is incident to a vertex of Z points locally into $(\mathbb{R}P^2)^-$.*

Proof. By statement of Lemma 2.8 we already know that $D_\otimes(Z)$ is a simple cycle and that each of its edges is even. By the construction of $D_\otimes$, the cycle $D_\otimes(Z)$ is homologically trivial on $\mathbb{R}P^2$. This implies that it splits $\mathbb{R}P^2$ into two components; see 'Background and notation'. For the rest, we need to inspect the construction of $D_\otimes$ in the proof of Lemma 2.8. However, we get all the required conclusions directly from the construction. □

Proofs of the lemmas. We are ready to present the proofs of Lemmas 2.18, 2.19 and 2.20. In all three proofs, $D_\otimes$ stands for the HT-drawing on $\mathbb{R}P^2$ from Lemma 2.36.

First, we consider Lemma 2.19, because its proof is the simplest one. In fact, we establish slightly stronger statement, which we will use again later. Lemma 2.19 follows directly from Lemma 2.37 below.

Lemma 2.37. *Let a, b, x and y be four distinct vertices of Z such that when we consider the two arcs of Z between a and b, the vertices x and y lie on the same arc. Then any two walks $\omega^+_{ab} \in W^+_{ab}$ and $\omega^+_{xy} \in W^+_{xy}$ must share a vertex.*

Proof. For contradiction, let $\omega^+_{ab} \in W^+_{ab}$ and $\omega^+_{xy} \in W^+_{xy}$ do not share a vertex. We consider a closed walk κ^+_{ab} arising from a concatenation of the walk ω^+_{ab} and the arc of Z connecting a and b not containing x, y. We also consider a closed walk κ^+_{xy} defined for x and y analogously to κ^+_{ab}; see Figure 2.17.

The closed curves $D_\otimes(\kappa^+_{ab})$ and $D_\otimes(\kappa^+_{xy})$ are both homologically non-trivial. Therefore, by Lemma 2.6, $D_\otimes(\kappa^+_{ab})$ and $D_\otimes(\kappa^+_{xy})$ must have an odd number of crossings. We note that either of the two curves may have self-intersections or self-touches, but there is a finite number of intersections between $D_\otimes(\kappa^+_{ab})$ and $D_\otimes(\kappa^+_{xy})$, which are necessarily crossings; see 'Background and notation'.

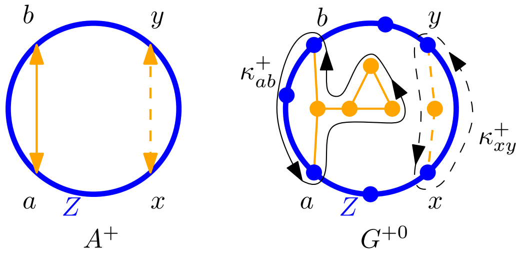

Figure 2.17: Walks in Lemma 2.37.

On the other hand, since $\omega^+_{ab} \in W^+_{ab}$ and $\omega^+_{xy} \in W^+_{xy}$ do not have a vertex in common, all pairs of edges in $E(\kappa^+_{ab}) \times E(\kappa^+_{xy})$ are independent. It follows that $D_\otimes(\kappa^+_{ab})$ and $D_\otimes(\kappa^+_{xy})$ have an even number of crossings, because $D_\otimes$ is an HT-drawing by Lemma 2.8; a contradiction. $\square$

We proceed to the proofs of Lemmas 2.18 and 2.20, which we restate here for reader's convenience.

Lemma 2.18. *Every inside arrow shares a vertex with every outside arrow.*

Proof. To the contrary, we assume that we have an inside arrow $\overline{xy}$ and an outside arrow $\overline{uv}$ which do not share any endpoint. Note that the arrows can be loops, that is, we allow $x = y$ or $u = v$. As before, we consider a closed walk κ^+_{xy} obtained from the concatenation of a walk from $\omega^+_{xy} \in W^+_{xy}$ and any of the two arcs of Z connecting x and y. If $x = y$, then we do not add the arc from Z. Analogously, we have a closed walk κ^-_{uv} coming from a walk in W^-_{uv} and an arc of Z connecting u and v. Both of these walks are non-trivial. We aim to contradict the intersection form on $\mathbb{R}P^2$ (see Lemma 2.6).

Unlike in the previous proof, this time $D_\otimes(\kappa^+_{xy})$ and $D_\otimes(\kappa^-_{uv})$ may not cross at every intersection. Namely, κ^+_{xy} and κ^-_{uv} may share a subpath of Z, but apart from the subpath the intersections are crossings; see Figure 2.18, left. To resolve the possible overlap, we slightly modify the drawing in the following way. Let us recall that $D_\otimes(Z)$ splits $\mathbb{R}P^2$ into two parts $(\mathbb{R}P^2)^+$ and $(\mathbb{R}P^2)^-$ according to Lemma 2.36. We push the subpath of κ^+_{xy} shared with Z, possibly consisting of a single vertex, a bit into $(\mathbb{R}P^2)^+$. This way, we obtain a drawing $D^+_\otimes$ of κ^+_{xy}. Similarly, we slightly push the subpath of κ^-_{uv} shared with Z into $(\mathbb{R}P^2)^-$; we denote the modified drawing of κ^-_{uv} by $D^-_\otimes$. The situation is depicted in Figure 2.18, right. Now $D^+_\otimes(\kappa^+_{xy})$ and $D^-_\otimes(\kappa^-_{uv})$ cross at every intersection and the crossings of $D^+_\otimes(\kappa^+_{xy})$ and $D^-_\otimes(\kappa^-_{uv})$ correspond to the crossings of $D_\otimes(\kappa^+_{xy})$ and $D_\otimes(\kappa^-_{uv})$.

Next, we consider the crossings of $D_\otimes(\kappa^+_{xy})$ and $D_\otimes(\kappa^-_{uv})$. Whenever e and f are two independent edges such that $e \in E(\kappa^+_{xy})$ and $f \in E(\kappa^+_{uv})$, then $D_\otimes(e)$ and $D_\otimes(f)$ have an even number of crossings, because $D_\otimes$ is an HT-drawing. Moreover, if e and f are adjacent, they cross evenly as well, since one of them must belong to Z. Here it is crucial that $\overline{xy}$ and $\overline{uv}$ do not share any endpoint. Therefore, $D_\otimes(\kappa^+_{xy})$ and $D_\otimes(\kappa^-_{uv})$ have an even number of crossings, and consequently, $D^+_\otimes(\kappa^+_{xy})$ and $D^-_\otimes(\kappa^-_{uv})$ as well. This is a contradiction to Lemma 2.6. $\square$

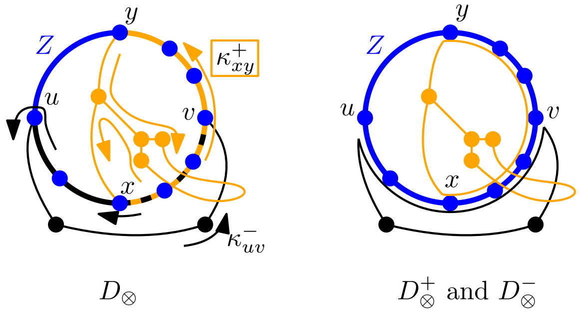

Figure 2.18: Walks in Lemma 2.18. The subpath of Z shared with both κ^+_{xy} and κ^-_{uv} is dashed in black and orange.

Lemma 2.20. *There are no three distinct vertices a, b, c on Z, an inside bridge B^+ and an outside bridge B^- such that B^+ induces the arrows $\overline{ab}$ and $\overline{ac}$ (and no other arrows) and B^- induces the arrows $\overline{ab}$ and $\overline{bc}$ (and no other arrows).*

Proof. For contradiction, there is such a configuration. Let e^+_a be an edge of $E(B^+)$ incident to a. Analogously, we define edges e^-_a, e^+_b, e^-_b, e^+_c and e^-_c.

Since B^+ cannot be just an edge, we have that $V(B^+) \setminus V(Z) \neq \emptyset$. Thus, by the definition of the bridges (see Definition 2.16), there is a proper walk in B^+ using the edges e^+_a and e^+_b; we denote it by ω^+_{ab}. This walk is non-trivial by Lemma 2.34. The assumptions of the lemma are satisfied by Proposition 2.33, since B^+ does not induce any inside loops.

Let κ^+_{ab} be the closed walk obtained from the concatenation of ω^+_{ab} and the arc of Z connecting a and b and avoiding c. Analogously, we define ω^+_{ac}, ω^-_{ab}, ω^-_{bc} and closed walks κ^+_{ac}, κ^-_{ab} and κ^-_{bc}. When defining the closed walks, we always use the arc of Z which avoids the third point among a, b and c. All the eight walks must be non-trivial by Lemma 2.34.

Now we show that e^+_a and e^-_a cross oddly in the drawing $D_\otimes$. For this, we consider the closed walks κ^-_{ab} and κ^+_{ac} and their drawings $D_\otimes(\kappa^-_{ab})$ and $D_\otimes(\kappa^+_{ac})$. The walks κ^-_{ab} and κ^+_{ac} share only the point a; therefore, $D_\otimes(\kappa^-_{ab})$ and $D_\otimes(\kappa^+_{ac})$ cross at every intersection, possibly except at $D_\otimes(a)$. By Lemma 2.36 we see that e^+_a and e^-_a point at a to different sides of Z in $D_\otimes$. Hence, $D_\otimes(\kappa^-_{ab})$ and $D_\otimes(\kappa^+_{ac})$ just touch in $D_\otimes(a)$, they cannot cross there. The touches can be removed by a slight perturbation of the drawings of the walks, similarly as in the proof of Lemma 2.18, without affecting any other intersections of the walks κ^-_{ab} and κ^+_{ac}. By Lemma 2.6 we therefore get that $D_\otimes(\kappa^-_{ab})$ and $D_\otimes(\kappa^+_{ac})$ have an odd number of crossings. However, considering any pair of edges (e, f) where $e \in E(\kappa^-_{ab})$ and $f \in E(\kappa^+_{ac})$ different from (e^-_a, e^+_a), we see that e and f cross an even number of times: indeed, for any such $(e, f) \neq (e^-_a, e^+_a)$, either e or f belongs to Z, or they are independent. Consequently, the odd number of crossings of $D_\otimes(\kappa^-_{ab})$ and $D_\otimes(\kappa^+_{ac})$ has to be realised on e^+_a and e^-_a.

Analogously, we show that e^+_b and e^-_b must cross oddly considering the walks κ^+_{ab} and κ^-_{bc}.

Now let us consider the closed walk κ^+_{ab} and a closed walk μ^-_{ab} obtained from the concatenation of ω^-_{ab} and the arc of Z connecting a and b which contains c. By the same reasoning as above, we get that $D_\otimes(\kappa^+_{ab})$ and $D_\otimes(\mu^-_{ab})$ touch in $D_\otimes(a)$ and $D_\otimes(b)$; if they intersect anywhere else, they cross there. Using a small perturbation as before, they must have an odd number of crossings by Lemma 2.6. On the other hand, the pairs of edges (e_a^+, e_a^-) and (e_b^+, e_b^-) cross oddly, as we have already observed. Any other pair (e, f) of edges where $e \in E(\kappa_{ab}^+)$ and $f \in E(\mu_{ab}^-)$ must cross evenly, since they are either independent, or one of them belongs to Z. In total, $D_\otimes(\kappa_{ab}^+)$ and $D_\otimes(\mu_{ab}^-)$ after the perturbation cross evenly, which is a contradiction. $\square$

Intersection of trivial interleaving walks. We conclude this section with a proof of a lemma similar in spirit to Lemma 2.37. We will need it in Section 2.6, but we keep the lemma here due to its similarity to previous statements.

Lemma 2.38. *Let a, b, x and y be four distinct vertices of Z such that x and y are on different arcs of Z when split by a and b. Let ω_{ab}^+ and ω_{xy}^+ be a proper ab-walk and a proper xy-walk in G^{+0}, respectively, such that $\lambda(\omega_{ab}^+) = \lambda(\omega_{xy}^+) = 0$. Then ω_{ab}^+ and ω_{xy}^+ must share a vertex.*

Proof. We proceed by contradiction. As usual, we consider closed walks κ_{ab}^+ and κ_{xy}^+ defined as follows. The walks κ_{ab}^+ consists of ω_{ab}^+ and an arc of Z connecting a and b, while the walk κ_{xy}^+ is formed by ω_{xy}^+ and an arc of Z connecting x and y. This time, ω_{ab}^+ and ω_{xy}^+ are trivial.

We push $D_\otimes(\kappa_{ab}^+)$ a bit inside and $D_\otimes(\kappa_{xy}^+)$ a bit outside of Z, similarly as in the proof of Lemma 2.18. This time, however, we introduce one more crossing, because both κ_{ab}^+ and κ_{xy}^+ are walks in G^{+0}. Lemma 2.6 implies that the drawings of κ_{ab}^+ and κ_{xy}^+ after the perturbation have to cross an even number of times. This, in turn, means that the drawings of ω_{ab}^+ and ω_{xy}^+ cross an odd number of times. Since $D_\otimes$ is an HT-drawing, it contradicts the assumption that ω_{ab}^+ and ω_{xy}^+ do not share a vertex. $\square$

2.6 Redrawing admissible configurations of arrows

In this section, we will prove Proposition 2.23 separately for each of the three cases. In other words, we show that if G^{+0} forms any of the configurations depicted in Figure 2.9, then G^{+0} admits an ordinary HT-drawing on S^2. Let us start with an auxiliary redrawing result that will be used in all three cases.

Lemma 2.39. *Let (D, λ) be a projective HT-drawing of G^{+0} on S^2 and Z be a cycle satisfying the separation assumptions. Let us also assume that $D(G^{+0}) \cap S^- = \emptyset$. Let B be one of the inside bridges different from an edge and let L be a valid labelling of B. Let us assume that there is at least one vertex $x \in V(B) \cap V(Z)$ such that $|L(x)| = 1$. Then there is a projective HT-drawing (D', λ') of G^{+0} on S^2 satisfying the following four conditions:*

(a) *D coincides with D' on Z and $D'(G^{+0}) \cap S^- = \emptyset$.*

(b) *Every edge $e \in E(G^{+0}) \setminus E(B)$ satisfies $\lambda(e) = \lambda'(e)$.*

(c) *Every edge $e \in E(B)$ that is not incident to Z satisfies $\lambda'(e) = 0$.*

(d) *For every edge $uv = e \in E(B)$ such that $u \in V(Z)$, we have $\lambda'(e) \in L(u)$.*

Note that the condition (b) allows that the edges in inside bridges other than B may be redrawn, but only under the condition that their triviality/non-triviality is not affected.

Proof. Let B^+ be the subgraph of B induced by the vertices of $V(B) \setminus V(Z)$. By the definition of a bridge (see Definition 2.16), the graph B^+ is connected. It is also non-empty, since we assume that B is not an edge.

Every cycle in the graph B^+ must be trivial. Indeed, if B^+ contained a non-trivial cycle, it could be used to obtain a proper, non-trivial walk from x to x. However, using Proposition 2.33, it would contradict the fact that $|L(x)| = 1$. That is, B^+ satisfies the assumptions of the planarization lemma (see Lemma 2.12). Let $U \subseteq V(B^+)$ be the set of vertices from the conclusion of Lemma 2.12. Consequently, if we perform the vertex-crosscap switches on U, we obtain a projective HT-drawing (D_U, λ_U) such that $\lambda_U(e) = 0$ for every edge $e \in E(B^+)$.

Let us recall that every vertex-crosscap switch over a vertex y is obtained by vertex-edge switches of non-trivial edges over y and then from swapping the value of λ on all edges incident to y. The vertex-edge switches do not affect the value of λ. Overall, we get that D_U coincides with D on Z. We also require that all vertex-edge switches are performed in S^+; therefore, D_U avoids S^-. Altogether, D_U and λ_U satisfy (a), (b) and (c), but we have not verified yet that (d) is satisfied as well.

In fact, (d) may not be satisfied at the moment and we may need to modify (D_U, λ_U). Let e_0 be any edge incident with x. If $L(x) = \{\lambda_U(e_0)\}$, we set $D' := D_U$ and $\lambda' := \lambda_U$. If $L(x) \neq \{\lambda_U(e_0)\}$, we perform vertex-crosscap switches over all vertices in $V(B^+)$, which yields D' and λ'. We have to check that (a)–(d) hold for D' and λ'.

It is sufficient to verify (a), (b) and (c) only in the latter case. Regarding (a), we have changed the drawing only by vertex-edge switches over edges e with $\lambda_U(e) = 1$ inside S^+. Validity of (b) is obvious from the fact that λ_U may have been changed only on edges incident with $V(B^+)$. Moreover, note that for every edge $e \in E(B^+)$ we have performed the vertex-crosscap switch for both endpoints of e. Therefore, $\lambda'(e) = \lambda_U(e) = 0$, which establishes (c). It remains to check (d).

First, we realise that we have set up D' and λ' in such a way that $L(x) = \{\lambda'(e_0)\}$—if $L(x) \neq \{\lambda_U(e_0)\}$, then we have made a vertex-crosscap switch over exactly one endpoint of e_0. In particular, we have just checked (d) for $e = e_0$.

To finish the proof, let $e = uv$ be an edge from the statement of (d) such that $e \neq e_0$. We need to check that $\lambda'(e) \subseteq L(u)$. This is trivially true if $L(u) = \{0, 1\}$, and thus, we may assume that $|L(u)| = 1$. Let ω be any proper xu-walk in B containing e_0 and e. The existence of such a walk is assured by the definition of a bridge (see Definition 2.16). We have $\lambda(\omega) = \lambda'(\omega)$, because the vertex-crosscap switches over the inner vertices of ω do not affect the triviality of ω. On the other hand, we can also write that $\lambda'(\omega) = \lambda'(e_0) + \lambda'(e)$, because $\lambda'(f) = 0$ for every edge $f \in E(B^+)$. Since $L(x) = \{\lambda'(e_0)\}$ and $|L(u)| = 1$, it follows that $L(u) = \{\lambda'(e)\}$ by Proposition 2.33 and Lemma 2.34 applied to x and u. □

For reader's convenience, we restate Proposition 2.23 here:

Proposition 2.23. *Let (D, λ) be a projective HT-drawing of G^{+0} on S^2 and Z be a cycle satisfying the separation assumptions. Moreover, let us assume that $D(G^{+0}) \cap S^- = \emptyset$ (that is, G^{+0} is fully drawn on $S^+ \cup D(Z)$). Let us also assume that G^{+0} forms an inside fan, an inside square or an inside split triangle. Then there is an ordinary HT-drawing D' of G^{+0} on S^2 such that D coincides with D' on Z and $D'(G^{+0}) \cap S^- = \emptyset$.*

Inside fan. We are ready to prove Proposition 2.23 for an inside fan, which is the simplest case. The idea behind the proof is to gradually alter the drawing in such a way that all non-trivial edges become incident to a single vertex and, at the same time, the new drawing will still be a projective HT-drawing. Once this is ensured, the value of λ may be simply set to zero for every edge yielding an ordinary HT-drawing.

Proof of Proposition 2.23 for an inside fan. We assume that G^{+0} forms an inside fan; see Figure 2.9. Let $x \in V(Z)$ be the endpoint common to all inside arrows. Let us consider an inside bridge B, possibly trivial[18]. Let $L = L_B$ be a valid labelling of B. It follows from Proposition 2.33 that $|L(u)| = 1$ for every $u \in V(B) \cap V(Z)$ different from x; there is at least one such u, since we assume that G is 2-connected[19]. In addition, all $u \in V(B) \cap V(Z)$ different from x must have the same label, because there are no arrows among them. Since we may switch all labels in a valid labelling by changing the value of the labelling parameter, we may assume that $L(u) = \{0\}$ for every $u \neq x$.

Next, we consider all inside bridges $B_1, \ldots, B_\ell$ and the corresponding labellings $L_{B_1}, \ldots L_{B_\ell}$ as above. We apply Lemma 2.39 to each of the bridges

[18]Recall that a bridge is trivial if it does not induce any arrow.

[19]This is contained in the separation assumptions; see Definition 2.15.

which are not an edge, one by one. This way we get a projective HT-drawing (D_1, λ_1), which satisfies:

(i) D coincides with D_1 on Z and $D_1(G^{+0}) \cap S^- = \emptyset$.

(ii) Every edge $e \in E(G^{+0})$ which is not incident with Z satisfies $\lambda_1(e) = 0$.

(iii) Every edge $e \in E(G^{+0})$ such that $\lambda_1(e) = 1$ is incident with x.

Property (i) is a consequence of the iterative application of property (a) from Lemma 2.39. Property (ii) follows from the iterative application of properties (b) and (c) of Lemma 2.39. Finally, property (iii) follows from (ii), from the iterative application of properties (b) and (d) of Lemma 2.39 and from the fact that every non-trivial inside bridge which is a single edge must contain x.

Finally, we set $D' := D_1$. Let $\lambda' \colon E(G^{+0}) \to \{0, 1\}$ be the constant function that assigns zero to every edge. We observe that from (ii) and (iii) it follows that $\lambda'(e)\lambda'(f) = \lambda_1(e)\lambda_1(f)$ for every pair of independent edges in $E(G^{+0})$. Therefore (D', λ') is a projective HT-drawing as well. Nevertheless, since λ' is identically zero, D' is also an ordinary HT-drawing on S^2. □

Inside square. Now we prove Proposition 2.23 for an inside square. Let B be the inside bridge inducing the inside square and let a, b, c and d be the vertices of $V(B) \cap V(Z)$ labelled according to Definition 2.21; see also Figure 2.9. The main ingredient for the present proof of Proposition 2.23 is the following lemma, which shows that B must have a suitable cut vertex.

Lemma 2.40. *The inside bridge B inducing the inside square contains a vertex v such that the graph $B - v$ is disconnected and the vertices a, b, c and d belong to four different components of $B - v$.*

We first show how to establish Proposition 2.23 for the case of an inside square using Lemma 2.40. The proof is analogous to the proof for an inside fan.

Proof of Proposition 2.23 for an inside square. We assume that B is the unique inside bridge inducing the inside square and a, b, c and d are vertices of $V(B) \cap V(Z)$ as above. In addition, let v be the cut vertex from Lemma 2.40.

First, we consider valid labellings of trivial inside bridges. Possibly after a switch of the value of the labelling parameter, we may ensure that all labels of a trivial inside bridge are zero. We apply Lemma 2.39 to all trivial inside bridges which are not an edge. We get a projective HT-drawing (D_1, λ_1) such that $\lambda_1(e) = 0$ for every edge $e \in E(G^{+0}) \setminus E(B)$. Note that we have not affected the values of λ on $E(B)$, D_1 coincides with D on Z and it still holds that $D_1(G^{+0}) \cap S^- = \emptyset$.

Second, we consider a valid labelling L of B. By Proposition 2.33, every vertex in $V(B) \cap V(Z)$ has just one label. It is easy to check that, up to switching all labels, we have $L(a) = L(c) = \{1\}$ and $L(b) = L(d) = \{0\}$. We apply Lemma 2.39 to B and the labelling L. We get a projective HT-drawing (D_2, λ_2) such that the only edges $e \in E(G^{+0})$ with $\lambda_2(e) = 1$ are edges of B incident to a or c.

Next, let C_a and C_c be the components of $B - v$ which contain a and c, respectively. We perform vertex-crosscap switches over all vertices of C_a and C_c except a, c and v. We perform the switches inside S^+, as usual. As a result, we get a projective HT-drawing (D_3, λ_3) such that the only edges $e \in E(G^{+0})$ with $\lambda_3(e) = 1$ are the edges of B incident to v.

Finally, we let $D' = D_3$ and we set $\lambda'(e) = 0$ for every edge e of G^{+0}. Analogously to the proof for the inside fans, $\lambda_3(e)\lambda_3(f) = \lambda'(e)\lambda'(f)$ for every pair of independent edges of G^{+0}. Therefore, (D', λ') is a projective HT-drawing on S^2 and D' is also an ordinary HT-drawing on S^2, as required. $\square$

It remains to prove Lemma 2.40 to conclude the case of an inside square. We present a certain separation lemma for a general graph and then we show that it implies Lemma 2.40.

Lemma 2.41. *Let G' be an arbitrary connected graph and $A = \{a_1, \dots, a_4\} \subseteq V(G')$ be a set of four distinct vertices. Let us assume that every a_ia_j-path has a common point in $V(G')\setminus A$ with every a_ka_ℓ-path whenever $\{i, j, k, \ell\} = \{1, 2, 3, 4\}$. Then there is a cut vertex v of G' such that $a_1, \dots, a_4$ are in four distinct components of $G' - v$.*

Proof. Let us consider an auxiliary graph G'' which is obtained from G' by adding two new vertices x, y and attaching x to a_1, a_2 and y to a_3, a_4. By the assumptions, G'' is connected, and moreover, there are no two vertex-disjoint paths connecting x and y. By Menger's theorem (see 'Background and notation'), there is a cut-vertex $v \in V(G'')\setminus\{x, y\} = V(G')$ disconnecting x and y. Let C_1 be the connected component of $G'' - v$ containing x and C_2 be the component containing y. Let C'_i, for $i = 1, 2$, be the subgraph of G' induced by v and the vertices of $C_i \cap G'$. Note that, since G' is connected, both C'_1 and C'_2 are connected. We show that v is the desired cut vertex.

Let p_1 be an a_1a_2-path in C'_1 and p_2 an a_3a_4-path in C'_2. Since C'_1 and C'_2 are connected, such paths p_1 and p_2 exist. Moreover, p_1 and p_2 may intersect only in v; however, according to the assumptions, they have to intersect in a vertex outside A. Therefore, they must intersect in v and $v \notin A$. Overall, we have verified that every a_ia_j-path passes through v, for $1 \le i < j \le 4$, which shows that v is the desired cut vertex. $\square$

Proof of Lemma 2.40. We apply Lemma 2.41 to B and to $A = \{a, b, c, d\}$. Let us consider a valid labelling L of B. Up to the swap of the labels, we may assume that $L(a) = L(c) = \{1\}$ and $L(b) = L(d) = \{0\}$. Then Proposition 2.33 together with Lemma 2.34 imply that every proper *ab*, *bc*, *cd* or *ad*-walk is nontrivial, whereas every proper *ac* or *bd*-walk is trivial. Then, the assumptions of Lemma 2.41 are satisfied due to Lemmas 2.37 and 2.38. $\square$

Inside split triangle. Finally, we prove Proposition 2.23 for an inside split triangle.

Proof of Proposition 2.23 for an inside split triangle. Let a, b, c be the three vertices of Z from the definition of the inside split triangle; see Definition 2.21 or Figure 2.9.

First, similarly as in the proof for the inside squares, we take care of trivial inside bridges via suitable labellings and Lemma 2.39. We obtain a projective HT-drawing (D_1, λ_1) that still satisfies the assumptions of Proposition 2.23, which,

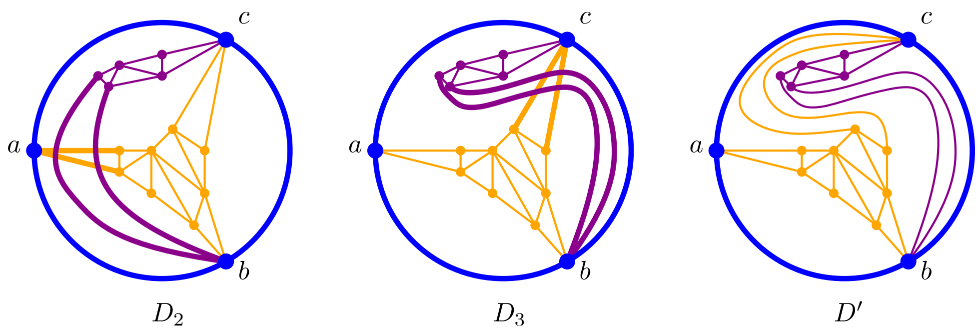

Figure 2.19: An example of redrawing an inside split triangle with one a-bridge and one bc-bridge. The edges participating in independent pairs crossing oddly are thick. For simplicity of the picture, the drawings D_3 and D' are simplified. For example, the vertex-edge switches used to obtain D_3 from D_2 introduce many pairs of independent edges crossing evenly and some pairs of adjacent edges crossing oddly. These intersections are removed in the picture as they do not play any role in the argument. In particular, the drawing D' is, in fact, typically not an embedding.

in addition, satisfies $\lambda_1(e) = 0$ for every edge $e \in E(G^{+0})$ that does not belong to a non-trivial bridge.

Now, let us consider non-trivial inside bridges. By the assumptions, every such bridge is either an *a-bridge*, that is, a non-trivial inside bridge which contains a (and b or c or both), or a *bc-bridge* which contains b and c, but not a. We consider valid labellings of these bridges. By Proposition 2.33, as before, a valid labelling assigns only one label to every vertex of the bridges lying on Z. As always, we may swap all labels in a valid labelling whenever needed. This way, it is easy to check that every a-bridge B admits a valid labelling L_B such that $L_B(a) = \{1\}$, whereas all other labels are 0. Similarly, each bc-bridge B admits a valid labelling L_B such that $L_B(b) = \{1\}$ and $L_B(c) = \{0\}$. We apply Lemma 2.39 to get a projective HT-drawing (D_2, λ_2) that still satisfies the assumptions of Proposition 2.23, which, in addition, has the following property: the edges $e \in E(G^{+0})$ with $\lambda_2(e) = 1$ are exactly the edges of an a-bridge which are incident to a, or edges of a bc-bridge incident to b.

If there are no bc-bridges, then all non-trivial edges are incident to a. We then finish the proof by setting $D' = D_2$ and letting λ' be identically equal to zero, similarly as in the cases of an inside fan and an inside square. However, if there are bc-bridge(s), we have to be more careful.

Let E_a^x and E_{bc}^x be the sets of edges incident to a vertex x in an a-bridge and the set of edges incident to x in a bc-bridge, respectively. Because D_2 is a projective HT-drawing, we have $\lambda_2(e)\lambda_2(f) = \mathrm{cr}_{D_2}(e, f)$ for every pair of independent edges e and f. In particular, $\mathrm{cr}_{D_2}(e, f) = 1$ for a pair of independent edges if and only if one of the edges belongs to E_a^a and the second one to E_{bc}^b.

Now, for every edge $e \in E_{bc}^b$, we perform the vertex-edge switch over each vertex different from a, b, c of every a-bridge. We call the resulting drawing D_3. The switches are performed inside S^+. This way, we change the crossing number of $e \in E_{bc}^b$ with the edges from E_a^a, E_a^b and E_a^c. In particular, after the redrawing, we get that $\mathrm{cr}_{D_3}(e, f) = 1$ for a pair of independent edges if and only if one of the edges belongs to E_a^c and the second one to E_{bc}^b. See Figure 2.19.

Finally, for every edge $e \in E_a^c$, we perform the vertex-edge switch over each vertex different from b and c of every bc-bridge obtaining the final drawing D'. Again, we perform the switches inside S^+. As a result, we change the crossing number of $e \in E_a^c$ with edges from E_{bc}^b and E_{bc}^c. However, it means that $\mathrm{cr}_{D'}(e, f) = 0$ for every pair of independent edges. That is, D' is the required ordinary HT-drawing on S^2. See Figure 2.19. $\square$

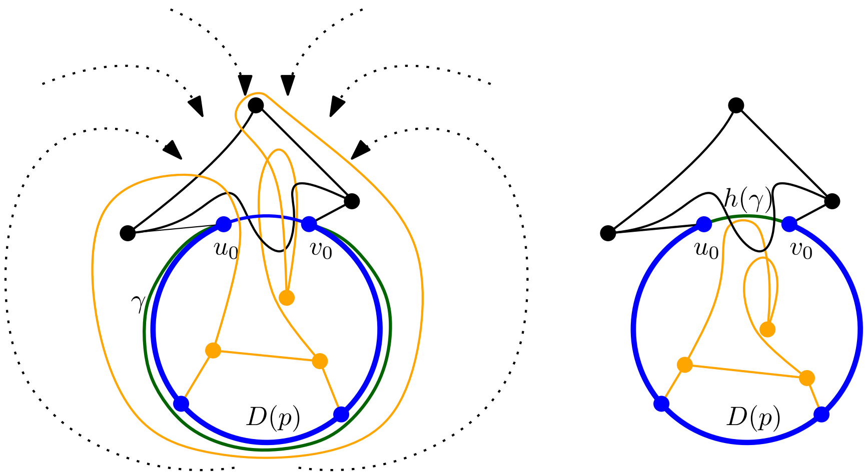

Figure 2.20: An illustration of the self-homeomorphism h, which maps B to S^+, applied to the drawing of $G^{+0} - e_0$ (where $e_0 = u_0v_0$).

2.7 Redrawing by Pelsmajer et al. [22]

It remains to prove Theorem 2.24. Again, we restate it here for convenience. As mentioned before, the present proof is almost identical to the proof by Pelsmajer et al. [22, Thm 2.1]. The only notable difference is that we avoid contractions[20]. As noted before, the proof by Fulek et al. [10, Lem. 3] can also be extended to yield the desired result.

Theorem 2.24. *Let D be a drawing of a graph G on the sphere S^2. Let Z be a cycle in G such that every edge of Z is even and Z is drawn as a simple cycle. Then there is a drawing D'' of G such that*

- *D'' coincides with D on Z,*
- *$D''(G^{+0})$ belongs to $S^+ \cup D(Z)$ and $D''(G^{-0})$ belongs to $S^- \cup D(Z)$*
- *and whenever (e, f) is a pair of edges such that both e and f are inside edges or both e and f are outside edges, then $\mathrm{cr}_{D''}(e, f) = \mathrm{cr}_D(e, f)$.*

Proof. First, we want to get a drawing such that there is only one edge of Z which may be intersected by some other edges. A part of the argument here is almost the same as the analogous argument in the proof of Lemma 2.28.

Let us consider an edge $e = uv \in E(Z)$ intersected by some other edges and let $f = vw \in E(Z)$ be an edge adjacent to e. We almost-contract e in the following way: we move the vertex v towards u until we remove all intersection of e with the other edges. This way, e becomes free of crossings and its former crossings appear now on f. Since both e and f were even edges in the initial drawing, f remains even after the redrawing as well. Also we do not affect parity of the other intersections and we remove possible self-intersections of the edges incident

[20]Contractions yield multi-graphs. Thus, we want to avoid contractions, because otherwise we would have to rework several notions for multi-graphs. Introducing multi-graphs in the previous sections would be disturbing and it is not convenient to repeat all the definitions in such a setting here.

with v similarly as in the proof of Lemma 2.28 (or see Figure 2.3, left). Finally, since we want to preserve the position of Z in the drawing, we consider a self-homeomorphism of S^2 which sends v back to its original position; see Figure 2.12.

With such redrawings it can be achieved that only one edge $e_0 = u_0v_0$ of Z may intersect the other edges while keeping Z fixed and e_0 even. Without loss of generality, we may assume that the original drawing D satisfies these assumptions.

Let p be the path in Z connecting u_0 and v_0 avoiding e_0. Let us also consider an arc γ connecting u_0 and v_0 outside (that is in S^-) close to $D(p)$ such that it does not cross any inside edge. The closed arc obtained from γ and $D(p)$ bounds two disks (2-dimensional balls). Let B be the open disk which contains S^+. Finally, we consider a self-homeomorphism h of S^2 that keeps $D(p)$ fixed and maps B to S^+. Considering the drawing $h \circ D$ on $G^{+0} - e_0$, it turns out that $G^{+0} - e_0$ is now drawn in S^+, up to p, which stays fixed. We also keep the original position of the edge e_0, that is, we do not apply h to this edge. See Figure 2.20.

Since the redrawing is done by a self-homeomorphism, we do not change the number of crossings among pairs of edges in G^{+0}. Analogously, we map G^{-0} to S^- and we get the required drawing. $\square$

Conclusion

In summary, we have presented an inductive procedure that gradually transforms a given HT-drawing on the projective plane to an embedding, and thus, provides a *constructive* proof of the strong Hanani–Tutte theorem on the projective plane. However, this is not the only interesting output. The tools that have been developed, most importantly the arrow graph, the labelling of the vertices of the bridges and the forbidden configurations of arrows expressed in Lemmas 2.18, 2.19 and 2.20 are of their own interest. To conclude the work, the author would like to comment on the advantages and disadvantages of these tools.

The arrows here were used to represent proper, non-trivial walks inside and outside. On other surfaces one would get not only one, but several types of arrows. The intersection form (see Lemma 2.6) becomes more complicated as well.

For instance, for the *Klein bottle*, which we can think of as a sphere with two crosscaps attached to it, we would get three non-trivial types of arrows according to the parity of crossings with the first crosscap and with the second crosscap. Lemma 2.18 then would say something like the following: every inside arrow of type α shares a vertex with every outside arrow of type β provided the intersection form guarantees an odd number of crossings between two closed curves of types α and β. Lemmas 2.19 and 2.20 can be generalised analogously.

We have used Lemmas 2.18, 2.19 and 2.20 to describe which configurations of arrows are admissible; in Proposition 2.22 we have shown that one of the three configurations from Definition 2.21 must appear either inside or outside. Similarly, one could use the generalised version of the lemmas to describe which configurations of arrows are admissible, say, on the Klein bottle. The author together with Martin Tancer and Pavel Paták have briefly thought in this direction and it seems that the number of configurations one would get for the Klein bottle with this approach counts roughly in tens.

On the one hand, the three configurations that were needed to be redrawn on the projective plane, as expressed in Proposition 2.22, provide a more compact description than 35 minimal forbidden minors for the projective plane. On the other hand, it seems that the number of admissible configurations of arrows on surfaces of higher genus is increasing with the genus, at least when one uses only the appropriate generalisations of the three forbidden configurations of arrows from Lemmas 2.18, 2.19 and 2.20. However, it is possible that there are additional topological restrictions not covered by the tools presented here.

The author would like to remark that the problem with the growing number of admissible configurations of arrows is not the only obstacle for the generalisation of the present approach to surfaces of higher genus. The other parts would have to be generalised as well and it is not clear that a sufficient generalisation is possible, even for the Klein bottle or the torus. Here, for instance, once we have fixed a trivial cycle Z which splits the graph and the surface into two parts, we were always able to redraw one of the sides without using the crosscap, and moreover, without changing the homological type of the cycle Z. On other surfaces, it may happen that it will not always be possible to preserve the type of the separating cycle Z during the redrawing, which would bring in additional difficulties when setting up the main induction.

On the positive side, in the light of the counter-example for the strong Hanani–

Tutte theorem on orientable surfaces of genus four and higher by Fulek and Kynčl [9], we can *speculate* that the theorem fails also for non-orientable surfaces of genus at least a certain small number. Then an approach based on ideas presented here may possibly be used, perhaps with an aid of a computer, to establish the remaining cases of small genus.

References

References for Chapter 1

[1] R. Baire. Sur les fonctions de variables réelles. *Annali di Matematica Pura ed Applicata (1898-1922)*, 3(1):1–123, 1899. ISSN 0373-3114.

[2] S. Bates, W. B. Johnson, J. Lindenstrauss, D. Preiss, and G. Schechtman. Affine Approximation of Lipschitz Functions and Nonlinear Quotients. *Geometric and Functional Analysis*, 9(6):1092–1127, 1999. ISSN 1420-8970. doi: 10.1007/s000390050108.

[3] Y. Benyamini and J. Lindenstrauss. *Geometric nonlinear functional analysis*, volume 48. American Mathematical Soc., 1998. ISBN 978-0-8218-0835-1.

[4] G. Blache, M. Karpinski, and J. Wirtgen. On approximation intractability of the bandwidth problem. *Electronic Colloquium on Computational Complexity (ECCC)*, 1998. ISSN 1433-8092. Technical report no. 14.

[5] M. Bonk and B. Kleiner. Rigidity for Quasi-Möbius Group Actions. *J. Differential Geom.*, 61(1):81–106, 05 2002. ISSN 0022-040X. doi: 10.4310/jdg/1090351321.

[6] J. Bourgain. On lipschitz embedding of finite metric spaces in Hilbert space. *Israel Journal of Mathematics*, 52(1):46–52, Mar 1985. ISSN 1565-8511. doi: 10.1007/BF02776078.

[7] L. E. J. Brouwer. Beweis der Invarianz des n-dimensionalen Gebiets. *Mathematische Annalen*, 71(3):305–313, 1911. ISSN 0025-5831.

[8] D. Burago and B. Kleiner. Separated nets in Euclidean space and Jacobians of biLipschitz maps. *Geometric and Functional Analysis*, 8:273–282, 1998. ISSN 1016-443X. doi: 10.1007/s000390050056.

[9] P. Z. Chinn, J. Chvátalová, A. K. Dewdney, and N. E. Gibbs. The bandwidth problem for graphs and matrices—a survey. *Journal of Graph Theory*, 6(3): 223–254, 1982. ISSN 0364-9024. doi: 10.1002/jgt.3190060302.

[10] M. I. Cortez and A. Navas. Some examples of repetitive, nonrectifiable Delone sets. *Geometry & Topology*, 20(4):1909–1939, 2016. ISSN 1364-0380.

[11] B. Dacorogna and J. Moser. On a partial differential equation involving the Jacobian determinant. *Annales de l'institut Henri Poincaré (C) Analyse non linéaire*, 7(1):1–26, 1990. ISSN 0294-1449. URL `http://eudml.org/doc/78211`.

[12] G. David. Morceaux de graphes lipschitziens et integrales singulières sur une surface. *Revista Matemática Iberoamericana*, 4(1):73–114, 1988. ISSN 02132230. URL `http://eudml.org/doc/39350`.

[13] G. David. Opérateurs d'intégrale singulière sur les surfaces régulières. *Annales scientifiques de l'École Normale Supérieure*, 21:225–258, 1988. ISSN 0012-9593.

[14] G. David and S. Semmes. *Fractured Fractals and Broken Dreams: Self-similar Geometry Through Metric and Measure*. Oxford lecture series in mathematics and its applications. Clarendon Press, 1997. ISBN 9780198501664.

[15] G. David and S. Semmes. Regular mappings between dimensions. *Publicacions Matemàtiques*, 44(2):369–417, 2000. ISSN 02141493, 20144350. URL `http://www.jstor.org/stable/43736696`.

[16] K. Deimling. *Nonlinear functional analysis*. Springer-Verlag, 1985. ISBN 978-3-662-00549-1. doi: 10.1007/978-3-662-00547-7.

[17] J. Díaz, J. Petit, and M. Serna. A survey of graph layout problems. *ACM Comput. Surv.*, 34(3):313–356, September 2002. ISSN 0360-0300. doi: 10.1145/568522.568523.

[18] C. Dubey, U. Feige, and W. Unger. Hardness results for approximating the bandwidth. *Journal of Computer and System Sciences*, 77(1):62–90, 2011. ISSN 0022-0000. doi: 10.1016/j.jcss.2010.06.006. Celebrating Karp's Kyoto Prize.

[19] N. Dunford and J. T. Schwartz. *Linear Operators: General theory, volume 1*, volume 7 of *Pure and applied mathematics*. Interscience Publishers, 1958. ISBN 978-0471608486.

[20] M. Dymond, V. Kaluža, and E. Kopecká. Mapping n grid points onto a square forces an arbitrarily large Lipschitz constant. *ArXiv e-prints*, April 2017. URL `https://arxiv.org/abs/1704.01940`.

[21] M. Dymond, V. Kaluža, and E. Kopecká. Mapping n grid points onto a square forces an arbitrarily large lipschitz constant. *Geometric and Functional Analysis*, Apr 2018. ISSN 1420-8970. doi: 10.1007/s00039-018-0445-z.

[22] H. Federer. *Geometric Measure Theory*. Classics in Mathematics. Springer, January 1996. ISBN 3540606564.

[23] U. Feige. Approximating the Bandwidth via Volume Respecting Embeddings. *Journal of Computer and System Sciences*, 60(3):510 – 539, 2000. ISSN 0022-0000. doi: 10.1006/jcss.1999.1682.

[24] I. Fonseca and W. Gangbo. *Degree Theory in Analysis and Applications*. Oxford Lecture Series in Mathematics and its Applications. Clarendon Press, 1995. ISBN 9780198511960.

[25] D. H. Fremlin. *Measure Theory*. Number vol. 2: Broad Foundations in Measure Theory. Torres Fremlin, 2000. ISBN 0953812928.

[26] A. I. Garber. On equivalence classes of separated nets. *Modelirovanie i Analiz Informatsionnykh Sistem*, 16(2):109–118, 2009. ISSN 1818-1015.

[27] M. L. Gromov. *Geometric Group Theory: Asymptotic invariants of infinite groups*. London Mathematical Society lecture note series. Cambridge University Press, 1993. ISBN 0521446805.

[28] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002. ISBN 9780521795401.

[29] W. B. Johnson, J. Lindenstrauss, D. Preiss, and G. Schechtman. Uniform quotient mappings of the plane. *Michigan Math. J.*, 47(1):15–31, 05 2000. ISSN 0026-2285. doi: 10.1307/mmj/1030374666.

[30] P. W. Jones. Lipschitz and bi-Lipschitz Functions. *Revista Matemática Iberoamericana*, 4(1):115–121, 1988. ISSN 0213-2230. URL `http://dmle.cindoc.csic.es/pdf/MATEMATICAIBEROAMERICANA_1988_04_01_06.pdf`.

[31] S. H. Jones. Applications of the Baire Category Theorem. *Real Anal. Exchange*, 23(2):363–394, 1999. ISSN 0147-1937.

[32] V. Kaluža. Lipschitz mappings of discrete sets. Bachelor thesis, Faculty of Mathematics and Physics, Charles University in Prague, 2012. [in Czech].

[33] V. Kaluža. Lipschitz mappings in the plane. Master thesis, Faculty of Mathematics and Physics, Charles University in Prague, 2014. `https://is.cuni.cz/webapps/zzp/detail/141371/?lang=en`.

[34] M. Kirszbraun. Über die zusammenziehende und Lipschitzsche Transformationen. *Fundamenta Mathematicae*, 22(1):77–108, 1934. ISSN 0016-2736. URL `http://eudml.org/doc/212681`.

[35] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, Jun 1995. ISSN 1439-6912. doi: 10.1007/BF01200757.

[36] L. O. Mafteiu-Scai. The Bandwidths of a Matrix. A Survey of Algorithms. *Annals of West University of Timisoara - Mathematics and Computer Science*, 52:183, 2015. ISSN 18413307. doi: 10.2478/awutm-2014-0019.

[37] A. N. Magazinov. The family of bi-Lipschitz classes of Delone sets in Euclidean space has the cardinality of the continuum. *Proceedings of the Steklov Institute of Mathematics*, 275(1):78–89, 2011. ISSN 1531-8605. doi: 10.1134/S0081543811080050.

[38] O. Maleva. Point Preimages under Ball Non-Collapsing Mappings. In *Geometric Aspects of Functional Analysis: Israel Seminar 2001-2002*, pages 148–157. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003. ISBN 978-3-540-36428-3. doi: 10.1007/978-3-540-36428-3_13.

[39] J. Matoušek and A. Naor (eds.). Open problems on low-distortion embeddings of finite metric spaces, 2011 (last revision). Available at `kam.mff.cuni.cz/~matousek/metrop.ps`.

[40] J. Matoušek (ed.). Open problems, workshop on discrete metric spaces and their algorithmic applications. KAM Series (Tech. report), Department of Applied Mathematics, Charles University, Prague, 2002. Available at `kam.mff.cuni.cz/~matousek/haifaop.ps`.

[41] P. Mattila. *Geometry of Sets and Measures in Euclidean Spaces: Fractals and Rectifiability*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1999. ISBN 9780521655958.

[42] C. T. McMullen. Lipschitz maps and nets in Euclidean space. *Geometric and Functional Analysis*, 8:304–314, 1998. ISSN 1016-443X. doi: 10.1007/s000390050058.

[43] J. R. Munkres. *Topology*. Featured Titles for Topology Series. Prentice Hall, Incorporated, 2000. ISBN 9780131816299.

[44] Ch. H. Papadimitriou. The NP-Completeness of the bandwidth minimization problem. *Computing*, 16(3):263–270, 1976. ISSN 1436-5057. doi: 10.1007/BF02280884.

[45] T. Rivière and D. Ye. Resolutions of the prescribed volume form equation. *Nonlinear Differential Equations and Applications NoDEA*, 3(3):323–369, 1996. ISSN 1021-9722. doi: 10.1007/BF01194070.

[46] W. Unger. The complexity of the approximation of the bandwidth problem. In *Proceedings 39th Annual Symposium on Foundations of Computer Science (Cat. No.98CB36280)*, pages 82–91, Nov 1998. ISBN 0-8186-9172-7. doi: 10.1109/SFCS.1998.743431.

[47] R. Viera. Densities non-realizable as the Jacobian of a 2-dimensional bi-Lipschitz map are generic. *arXiv preprint arXiv:1603.07310*, 2016. URL `https://arxiv.org/abs/1603.07310`.

[48] L. Zajíček. Porosity and σ-porosity. *Real Analysis Exchange*, 13(2):314–350, 1987. ISSN 01471937, 19301219.

[49] L. Zajíček. On σ-porous sets in abstract spaces. *Abstr. Appl. Anal.*, 2005(5):509–534, 2005. ISSN 1085-3375. doi: 10.1155/AAA.2005.509.

References for Chapter 2

[1] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *The design and analysis of computer algorithms*. Addison-Wesley series in computer science and information processing. Addison-Wesley Pub. Co., 1974. ISBN 978-0201000290.

[2] D. Archdeacon. A Kuratowski theorem for the projective plane. *Journal of Graph Theory*, 5(3):243–246, 1981. doi: 10.1002/jgt.3190050305.

[3] J. Battle, F. Harary, Y. Kodama, and J. W. T. Youngs. Additivity of the genus of a graph. *Bull. Amer. Math. Soc.*, 68:565–568, 1962. ISSN 0002-9904. doi: 10.1090/S0002-9904-1962-10847-7.

[4] G. Cairns and Y. Nikolayevsky. Bounds for Generalized Thrackles. *Discrete & Computational Geometry*, 23(2):191–206, 2000. ISSN 1432-0444. doi: 10.1007/PL00009495.

[5] Ch. Chojnacki. Über wesentlich unplättbare Kurven im dreidimensionalen Raume. *Fundamenta Mathematicae*, 23(1):135–142, 1934. URL `http://eudml.org/doc/212715`.

[6] É. Colin de Verdière, V. Kaluža, P. Paták, Z. Patáková, and M. Tancer. A Direct Proof of the Strong Hanani–Tutte Theorem on the Projective Plane. *Journal of Graph Algorithms and Applications*, 21(5):939–981, 2017. doi: 10.7155/jgaa.00445.

[7] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Heidelberg, fourth edition, 2010.

[8] D. B. Fuchs and O. Ya. Viro. Homology and cohomology. In V. A Rokhlin and S. P. Novikov, editors, *Topology II: Homotopy and Homology. Classical Manifolds*, pages 95–196. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004. ISBN 978-3-662-10581-8. doi: 10.1007/978-3-662-10581-8_2.

[9] R. Fulek and J. Kynčl. Counterexample to an extension of the Hanani–Tutte theorem on the surface of genus 4. *ArXiv e-prints*, September 2017. URL `https://arxiv.org/abs/1709.00508`.

[10] R. Fulek, M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Adjacent Crossings Do Matter. *Journal of Graph Algorithms and Applications*, 16(3):759–782, 2012. doi: 10.7155/jgaa.00266.

[11] R. Fulek, J. Kynčl, I. Malinović, and D. Pálvölgyi. Clustered planarity testing revisited. *The Electronic Journal of Combinatorics*, 22(4):P4–24, 2015.

[12] A. Gagarin, W. Myrvold, and J. Chambers. Forbidden minors and subdivisions for toroidal graphs with no K3, 3's. *Electronic Notes in Discrete Mathematics*, 22:151–156, 2005.

[13] H. H. Glover, J. P. Huneke, and C. S. Wang. 103 Graphs that are irreducible for the projective plane. *Journal of Combinatorial Theory, Series B*, 27 (3):332 – 370, 1979. ISSN 0095-8956. doi: doi.org/10.1016/0095-8956(79) 90022-4.

[14] A. Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002. ISBN 0-521-79160-X; 0-521-79540-0.

[15] K. Kawarabayashi, B. Mohar, and B. Reed. A Simpler Linear Time Algorithm for Embedding Graphs into an Arbitrary Surface and the Genus of Graphs of Bounded Tree-Width. In *49th Annual IEEE Symposium on Foundations of Computer Science, 2008.*, pages 771–780, Oct 2008. doi: 10.1109/FOCS.2008.53.

[16] C. Kuratowski. Sur le problème des courbes gauches en Topologie. *Fundamenta Mathematicae*, 15(1):271–283, 1930. URL `http://eudml.org/doc/212352`.

[17] R. B. Levow. On Tutte's algebraic approach to the theory of crossing numbers. In *Proceedings of the Third Southeastern Conference on Combinatorics, Graph Theory, and Computing*, pages 315–314. Florida Atlantic Univ., Boca Raton, Fla., 1972.

[18] G. L. Miller. An additivity theorem for the genus of a graph. *Journal of Combinatorial Theory, Series B*, 43(1):25 – 47, 1987. ISSN 0095-8956. doi: doi.org/10.1016/0095-8956(87)90028-1.

[19] B. Mohar. A Linear Time Algorithm for Embedding Graphs in an Arbitrary Surface. *SIAM Journal on Discrete Mathematics*, 12(1):6–26, 1999. doi: 10.1137/S089548019529248X.

[20] B. Mohar and C. Thomassen. *Graphs on Surfaces*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, 2001.

[21] J. R. Munkres. *Topology*. Prentice Hall, 2000.

[22] M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Removing even crossings. *Journal of Combinatorial Theory, Series B*, 97(4):489–500, 2007. doi: 10.1016/j.jctb.2006.08.001.

[23] M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Removing even crossings on surfaces. *Electronic Notes in Discrete Mathematics*, 29:85 – 90, 2007. ISSN 1571-0653. doi: 10.1016/j.endm.2007.07.015. European Conference on Combinatorics, Graph Theory and Applications.

[24] M. J. Pelsmajer, M. Schaefer, and D. Stasi. Strong Hanani–Tutte on the Projective Plane. *SIAM Journal on Discrete Mathematics*, 23(3):1317–1323, 2009. doi: 10.1137/08072485X.

[25] N. Robertson and P. D. Seymour. Graph minors. viii. a Kuratowski theorem for general surfaces. *Journal of Combinatorial Theory, Series B*, 48(2):255 – 288, 1990. ISSN 0095-8956. doi: doi.org/10.1016/0095-8956(90)90121-F.

[26] M. Schaefer. Hanani–Tutte and Related Results. In I. Bárány, K. J. Böröczky, G. F. Tóth, and J. Pach, editors, *Geometry — Intuitive, Discrete, and Convex: A Tribute to László Fejes Tóth*, pages 259–299. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013. ISBN 978-3-642-41498-5. doi: 10.1007/978-3-642-41498-5_10.

[27] M. Schaefer. Toward a Theory of Planarity: Hanani-Tutte and Planarity Variants. *Journal of Graph Algorithms and Applications*, 17(4):367–440, 2013. doi: 10.7155/jgaa.00298.

[28] M. Schaefer and D. Štefankovič. Block Additivity of $\mathbb{Z}_2$-Embeddings. In S. Wismath and A. Wolff, editors, *Graph Drawing: 21st International Symposium, GD 2013, Bordeaux, France, September 23-25, 2013, Revised Selected Papers*, pages 185–195, Cham, 2013. Springer International Publishing. ISBN 978-3-319-03841-4. doi: 10.1007/978-3-319-03841-4_17.

[29] A. Shapiro. Obstructions to the imbedding of a complex in a Euclidean space. I. The first obstruction. *Annals of Mathematics, Second Series*, 66: 256–269, 1957. doi: 10.2307/1969998.

[30] W. T. Tutte. Toward a theory of crossing numbers. *Journal of Combinatorial Theory*, 8(1):45 – 53, 1970. ISSN 0021-9800. doi: http://dx.doi.org/10.1016/S0021-9800(70)80007-2.

[31] E. R. van Kampen. Komplexe in euklidischen Räumen. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 9(1):72–78, 1933. ISSN 1865-8784. doi: 10.1007/BF02940628.

[32] W. Wu. On the realization of complexes in Euclidean spaces. I. *Acta Mathematica Sinica*, 5:505–552, 1955.