Relaxing and lifting triangulations

Doctoral Thesis

at

Graz University of Technology

submitted by

Thomas Hackl

April 12, 2010

Institute for Software Technology
and
Institute for Theoretical Computer Science,

Faculty of Computer Science,
Graz University of Technology
A–8010 Graz, Austria

Advisers: Dipl.-Ing. Dr.techn. Univ.Doz. Oswin Aichholzer
Dipl.-Ing. Dr.techn. Univ.-Prof. Franz Aurenhammer

Preface

Why? ... Why writing a thesis, this thesis, you ask? Well, I asked myself the same question. And the surprisingly simple answer is: To get a degree. (Note: Do not confuse with vertex degree or degree of a polygonal complex.) So why should you read it then, you ask? Now, unfortunately there's no simple answer to that question. You might be one of my two supervisors: well cling together, swing together. Maybe you are a reviewer: thank you very much for the effort and volunteering, and sorry you had to read through all of this. Others: confess, you're some kind of masochist. Don't deny it. You are even reading this preface, aren't you?

But what's the story behind me satisfying your masochistic tendencies? Well, everything started in the year 1975. If my parents only had known. Nine months later, on a Sunday, I came to this world, just in time for lunch. The first nineteen years went quite normal and predictable with the slight exception that I went to a Higher Technical School (HTL) to get my A-levels. After an eight months long diet from all which is logical and sane (call it national service), I thirsted for knowledge. So I came to Graz University of Technology. And now I am sitting here, writing on a preface that hardly anyone will ever read.

But lets make one step back. You start your journey to PhD-hood not really knowing what lies in front of you. You have just got your Masters degree and actually have no clue. (You know what you don't know, but don't know what you know.) The only thing you know is that there is THE THESIS that has to be written. But that will happen in a time far, far in the future. Then you meet very smart, interesting people that know of lots of very nice and interesting problems and questions. So you start working with them, getting results, and you publish some papers, first one, then two, and more following. You fly forth and back, visit fellow researchers and be visited by them, co-organize some nice workshops, and more papers get published. So time goes by and suddenly everybody starts asking the same stupid question: "When will you get your degree?" That is when you realize that almost five years have passed, and THE THESIS is a very complex thing, completely imaginary.

So you have a look at your publications. Most just do not fit the topic of your work (lots of interesting problems arose during the years and although they were from different other fields of research, I'm not here to judge, they equally have the need to be solved). But you realize with some satisfaction, that some of them can be used to build upon THE THESIS.

Of course, you were foresighted enough to always use (at least slightly) different notations in your papers. Because otherwise you really would have to few work to do. So notations and definitions have to be checked and adapted to fit together and not to be completely contradictory. You think about the structure of your thesis and realize that all the new chapters need introductions. Of course, THE THESIS itself wants an introduction, too ... and a conclusion while we're at it. But while you are writing and printing and reading and writing, and so on ... THE THESIS grows, paragraph after paragraph, page after page. You finally finish the introduction, also manage to write a short abstract. Your Bibliography suddenly fills five pages and you even write a CV, so that everybody knows, who's to blame.

And then, one day, there it is: THE THESIS shiny new and fresh from the printer, still very complex but now with a big real part. Then you think of the people who will (have to) read THE THESIS and you decide to write one more page ... THE PREFACE.

Thomas

Acknowledgments

First and foremost, my best thanks go to my parents, without their support and understanding I would have never come this far. They are always there for me and give me stability and secureness.

I am indebted to my thesis supervisor Franz Aurenhammer, for introducing me to the topic, for his advises and suggestions, for the long discussions and for reviewing and proof-reading this thesis. Many tricky problems did we tackle and sealed their fate with elegant proofs.

I would like to thank my second thesis supervisor Oswin Aichholzer for his faith in me, to employ me for his project, and giving me plenty of rope in my choices of research topics. We had a lot of fun working on lots of different fields with many international researchers. Additionally, I learned a lot besides our main fields of research including the fine art of project leading diplomacy.

My thank goes to my colleagues, fellow researchers and co-authors. I very much enjoyed our gatherings/meetings at conferences and mutual research visits, where many new ideas were born, we also had a lot of fun, and more often than not many exquisite feasts. There were lots of interesting problems we tried ourselves on and many informative and fruitful discussions. Not to mention the nice papers we had. I've learned a lot from all of you and I hope that our cooperation will continue in the future.

Special thanks go to my colleagues at the Institute of Software Technology at Graz University of Technology with whom I spent the last five years of my professional life. We enjoyed countless coffees together, with or without some cake, having nice chats and refreshing laughs. And many thanks to the three very nice people I share my office with. Alex, Birgit, and Wolfgang, thank you for the friendly, enjoyable working environment, your readiness to help, the countless research discussions, the chit-chatting, and your friendship.

I am also not forgetting the nice girls from the Office of the Dean of the Faculty of Computer Science. Thank you very much for helping me through all the bureaucracy.

And as I already published some papers, a few among them leading to this thesis, I want to thank the many anonymous referees, who gave valuable comments and helped to improve both the results and the presentation of our work.

Abstract

The relaxation of triangulations to pseudo-triangulations is already well known. We investigate the differences between triangulations and pseudo-triangulations with respect to the optimality criterion of minimal total edge length. We show that, although (especially pointed) pseudo-triangulations have significantly less edges than triangulations, the minimum weight pseudo-triangulation can be a triangulation in general.

We study the flip graphs for pseudo-triangulations whose maximum vertex degree is bounded by some constant. For point sets in convex position, we prove that the flip graph of such pseudo-triangulations is connected if and only if the degree bound is larger than 6. We present an upper bound of $O(n^2)$ on the diameter of this flip graph and also discuss point sets in general position.

Finally we relax triangulations even beyond the concept of pseudo-triangulations and introduce the class of pre-triangulations. When considering liftings of triangulations in general polygonal domains and flipping therein, pre-triangulations arise naturally in three different contexts: When characterizing polygonal complexes that are liftable to three-space in a strong sense, in flip sequences for general liftable polygonal complexes, and as graphs of maximal locally convex functions.

Keywords: Computational Geometry, triangulation, pseudo-triangulation, pre-triangulation, optimality, flip operations, flip graph, surface realization, locally convex function, lifting, projectivity.

Deutsche Fassung:
Beschluss der Curricula-Kommission für Bachelor-, Master- und Diplomstudien vom 10.11.2008
Genehmigung des Senates am 1.12.2008

EIDESSTATTLICHE ERKLÄRUNG

Ich erkläre an Eides statt, dass ich die vorliegende Arbeit selbstständig verfasst, andere als die angegebenen Quellen/Hilfsmittel nicht benutzt, und die den benutzten Quellen wörtlich und inhaltlich entnommene Stellen als solche kenntlich gemacht habe.

Graz, am …………………………… ………………………………………………..
(Unterschrift)

Englische Fassung:

STATUTORY DECLARATION

I declare that I have authored this thesis independently, that I have not used other than the declared sources / resources, and that I have explicitly marked all material which has been quoted either literally or by content from the used sources.

…………………………… ………………………………………………..
date (signature)

Contents

List of Figures

Chapter 1

Introduction

Computational Geometry is a relatively young and very active field of research in computer science. Studying algorithms and data structures have been two main objectives of this ever-growing discipline [77] from the beginning. Today, Computational Geometry is often understood as "the study of geometric problems from a computational point of view" [46] and has become highly interweaved with fields of Discrete Geometry like Combinatorial Geometry. In this context, polygonal complexes soon became objects of interest from various points of view. Although they are geometric structures, they have a highly discrete nature. Straight lines spanned by a finite set of discrete points give rise to simple and memory efficient data structures. While not loosing the geometric information, polygonal complexes additionally provide combinatorial context that is sufficient for most applications and allows for very efficient and stable algorithms (see [28] for an introductory overview). For an example of this fact, see the Computational Geometry Algorithms Library (CGAL) [32].

Line arrangements arise as duals of finite point sets and, among many other applications, prove useful for efficiently finding degeneracies, like collinear points in the plane or coplanarities in spatial point sets [38]. Generalizations, such as arrangements of pseudo-lines, have already been well studied and turned out to be crucial for the enumeration of combinatorially different point sets of small size, that is, point configurations with different order type [9].

Another classical example for an important class of polygonal complexes are Voronoi diagrams, which feature proximity information between geometric objects [21, 43]. Thus they are the data structure of choice to solve the Post Office Problem, but are also prominent for their dual structure, the Delaunay triangulation, and their connection to convex polyhedra [39]. Generalizations of this versatile structure exist for higher-dimensional point sets [33, 60], other types of sites (like line segments instead of points) [28], different metrics [35, 65, 67], weighted distances (multiplicative or additive weights) [22, 42], or different distance functions like in so called Power diagrams [18, 19, 20, 24].

Although already very well studied, many new results are still made present-day. Most of these investigations deal with generalizations, like for instance the aspect-ratio (segment) Voronoi diagram [17]: A point p in the plane belongs to the Voronoi cell of the site s (segment) which maximizes the aspect-ratio for p, the ratio of the longest side to the smallest height of the triangle (p, s). Nevertheless, there also exist recent innovation on the classical Voronoi diagram, like a new approach of construction [3], that also works on sites consisting of straight lines and circular arcs, and allows for an efficient and robust implementation [16].

The probably most famous class of polygonal complexes are triangulations. Used in computer graphics for animations in movies and computer games, a wide range of people have already seen triangulations, but most probably not recognizing them as such. Triangulations partition complex geometric domains like, for instance, point sets, polygons, or chocolate disks (see Figure 1.1) into simple geometric objects (triangles, general simplices in higher dimensions, or just bite-sized pieces of delicious chocolate).

Figure 1.1: Pumpkin seed chocolate Mitzi Blue® [90] from Zotter [89]. Strictly speaking not a triangulation, as it is not even a polygonal complex (see Chapter 4), but still nice.

The before mentioned computer graphics is by far not the only field of application for triangulations. In the early days the development came from the need of cartographers and surveyors to calculate the elevation of certain points from already known points of reference. One famous example is the triangulation of the Kingdom of Hannover done by Carl Friedrich Gauß. This triangulation was on the backside of the 10 DM bill (Deutsche Mark), the front of this bill shows Gauß himself, see Figure 1.2. Since then, triangulations found many fields of application and are used in computer aided design (CAD) systems, as well as in finite element methods (FEM). Additionally, triangulations 'went back to the roots' and are employed in geographical information systems (GIS). The increasing amount of available data made the usage of computers inevitable, and thus the need for algorithms to construct accurate and correct terrain models [84]. And last but not least, triangulations are a valuable data structure in many algorithms and solutions in Computational Geometry itself, with the Delaunay triangulation as the most prominent representative. Not only does the Delaunay triangulation host a variety of very useful properties (like "nice" triangles, minimal angle maximized), and is a superset of many important geometric structures (like minimum spanning tree, Gabriel graph, relative neighborhood graph), but there also exist fast and reliable algorithms to construct this special triangulation [43, 44]. In addition, and as already mentioned, the duality between the Delaunay triangulation and the Voronoi diagram puts great importance to both structures.

Another widely known triangulation is the so called minimum weight triangulation, where the sum of all edge lengths is minimal over all possible triangulations (on the same set of vertices). Minimal edge length, which is one of the few interesting properties the Delaunay triangulation does not provide, turned out to be notoriously hard to achieve. Being a long standing open problem, Mulzer and Rote [70] only recently showed that calculating the minimum weight triangulation is actually an NP-hard problem.

Finally, an important feature of triangulations is that they can be transformed into each other

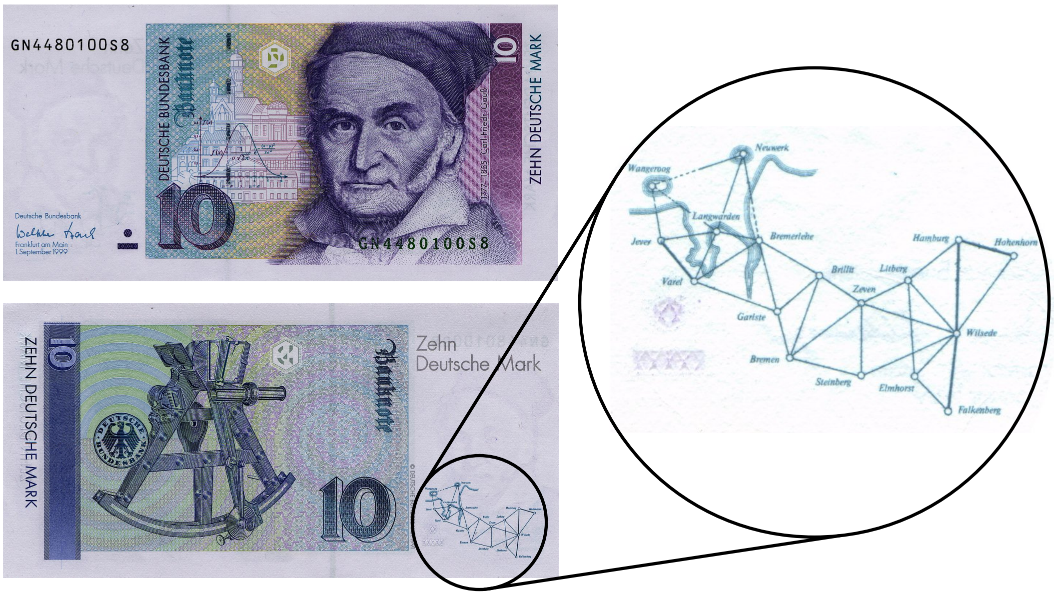

Figure 1.2: Front and back of the 10 Deutsche Mark bill [37] and a magnification of the triangulation of the Kingdom of Hannover done by Gauß.

by a series of local, constant sized transformations. These so called flips were first introduced by Lawson [63]. Therefore, flips on triangulations are often also called “Lawson flips”. Basically, a Lawson flip is executed in a convex quadrilateral that is a subset of the triangulation. The operation removes the diagonal of the quadrilateral that is part of the triangulation, and inserts the other diagonal of the quadrilateral. Thus a flip exchanges one edge of the triangulation with another one and is also often called an edge flip. For constructing Delaunay triangulations, applying an edge flip is restricted to edges that violate the local Delaunay property. Applying multiple transformations in this way leads to a directed sequence of flips, that always terminates in the Delaunay triangulation [61, 83]. Hence, this restricted version is often called a “Delaunay flip”.

Although triangulations (especially the Delaunay triangulation) have been extensively studied in the past, a relaxation has been discovered only in recent years. The so called pseudo-triangulations [79] were first introduced by Pocchiola and Vegter [74, 75, 76] as the maximal graph of crossing free bitangents between a set of disjoint convex objects in the plane. At more or less the same time, pseudo-triangulations appeared as geodesic triangulations [34, 48] in tilings of polygons using crossing free geodesics.

Since then, pseudo-triangulations gained importance in many fields of application (until then solely occupied by triangulations) expeditiously; among them being ray shooting [34, 48], visibility [72, 75], guarding [73, 85], rigidity [51, 78, 87, 88], and kinetic collision detection [2, 58, 59] to name only a few.

Pseudo-triangulations share many properties of triangulations like planarity, simple face shape, and the applicability of local flips. There exists even a new kind of edge flip [10, 71] for pseudo-triangulations, that allows for each interior edge to be flipped, which is not guaranteed for triangulations.

For triangulations on point sets, each extremal point has one angle larger than π between consecutive incident edges. Such points are often called “pointed”. In contrast to triangulations where every non-extremal point is not pointed, interior points may be pointed in pseudo-triangulations. If all points are pointed, then this is called a pointed pseudo-triangulation, or sometimes minimum pseudo-triangulation, because the number of edges is minimal in this case. Thus pointed pseudo-triangulations are a compact alternative for triangulations and allow more efficient algorithms, kinetic collision detection [58] being one example for the exploitation of this benefit.

1.1 Thesis overview

There exist many ways to triangulate a set of points in the plane [12]. For many applications (for example in computer graphics, where we don't want long and skinny triangles) we would like to find a triangulation with "nicely shaped" triangles, that is, triangles that are close to equilateral.

With this wish in mind, certain optimal triangulations were defined, where "optimal" depends on the chosen optimization criteria [26, 29]. For example, as equilateral triangles are also equiangular, we want the angles of a triangle to be more or less equal, or at least not too big or too small. In this respect, for maximizing the minimal occurring angle, the optimal triangulation is the Delaunay triangulation. As long, skinny triangles also have long edges, we can ask for minimizing the overall edge length. The optimal triangulation fulfilling this requirement is the minimum weight triangulation. In contrast to the very efficiently computable Delaunay triangulation, the complexity of the minimum weight triangulation was a long standing open problem [45], which has been shown to be NP-hard only recently by Mulzer and Rote [70].

Thus approximation algorithms for the minimum weight triangulation were studied intensively [26]. As pseudo-triangulations (in general) may have less edges than triangulations it is natural to ask for minimum weight pseudo-triangulations [81]. Gudmundsson and Levcopoulos [50] present two approximation algorithms for the minimum weight pseudo-triangulation of point sets. Further they show a cubic time algorithm to compute the minimum weight pseudo-triangulation of a simple polygon.

The question whether the minimum weight pointed pseudo-triangulation is always at least as good as the minimum weight triangulation remains open. Considering the fact that pointed pseudo-triangulations have a minimum number of edges over all pseudo-triangulations of a point set, one other question immediately arises: Is the minimum weight pseudo-triangulation always pointed? In Chapter 2 of this thesis we will investigate these questions. Further we will have a look at the so called greedy pseudo-triangulation, which is constructed by inserting edges in ascending order of their lengths.

After this warm-up, we will take a closer look at flipping in triangulations and pseudo-triangulations. It is well known that the flip graph of triangulations is connected [61, 83], that is, every triangulation of a point set can be reached from any other one with a sequence of edge flips. The diameter of the flip graph, and therefore the length of the flip sequence between two triangulations, is bounded with $O(n^2)$ [29, 56]. For pointed pseudo-triangulations, Bereg [27] proves a diameter of $O(n \log n)$ for the flip graph. Allowing the use of so called removing flips (and their inverse operation, the inserting flip) Aichholzer et al. show in [10] that every pseudo-triangulation (including triangulations) can be reached from every other pseudo-triangulation with at most $O(n \log n)$ edge flips.

Beyond that, also subgraphs of the flip graph of triangulations for certain constraints are of interest. Houle et al. [54] show that when permitting only triangulations that host a perfect matching, the flip graph is still connected. A similar result is obtained for order-k Delaunay graphs. Abellanas et al. [1] show that for point sets in convex position, order-k Delaunay graphs (see the introduction to Chapter 3 for a brief definition) are connected via edge flips for any $k > 0$. For point sets in general position, they prove that this is true for $k \leq 1$ and they show examples for $k > 1$ where the resulting flip graph is not connected.

Restricting flips to triangulations (and pseudo-triangulation) which do not violate a certain vertex degree, stems primarily from results on pseudo-triangulations in this context. The degree of a vertex is the number of edges incident to that vertex in a geometric graph, that is, a pseudo-triangulation in this case. Kettner et al. [57] show that, for every point set in the plane, there exists a pseudo-triangulation with a maximal vertex degree of 5. And Aichholzer et al. [14] extend this to constrained pseudo-triangulations and prove that also for every simple polygon there exists a pseudo-triangulation with maximal vertex degree of 5. These results are important ingredients for the kinetic data structure for collision detection presented in [58].

In this spirit we investigate the flip graph of vertex degree constrained pseudo-triangulations in Chapter 3. We prove connectedness of this flip graph for point sets in convex position. Here all pseudo-triangulations are actually triangulations. We also give a bound for the flip distance for this setting. For general point sets and polygons, we provide some negative examples for triangulations, showing that the flip graphs in those cases need not be connected.

Familiar with the idea of pseudo-triangulations and flipping therein, we will lift polygonal complexes to surfaces and derive generalizations for both concepts, in Chapter 4. It is easy to see that every triangulation of a point set in the plane can be easily lifted to a surface, by any arbitrary height vector applied to the vertex set. This is not so obvious for pseudo-triangulations. Aichholzer et al. [10] prove that this relaxation of triangulations can be lifted, when a certain, well defined subset of points is assigned a height vector, while the height for the remaining points is uniquely defined by this vector. The flip operation has also a representation in such surfaces. In liftings from triangulations, a so called surface flip exchanges a reflex edge with a convex edge, or vice versa. This is the analogue to the Lawson flip in two-dimensional triangulations. A similar flip exists in surfaces from pseudo-triangulations, with the addition of the edge-removing flip, which has kind of a “planarizing” behavior in surfaces.

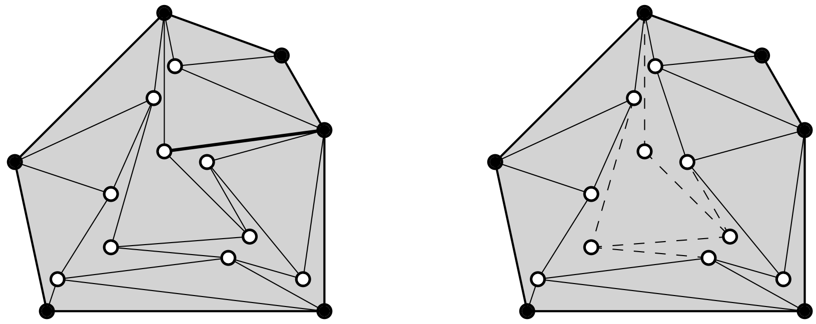

Figure 1.3: Flipping an edge (bold) in a pointed pseudo-triangulation (left) can result in the disappearance of edges (dashed) and vertices (solely incident to dashed edges) in every lifting of the resulting complex (right).

In contrast to pseudo-triangulations in the plane, when flipping in surfaces from pseudo-triangulations several edges may disappear after a flip. This happens when neighbored faces in different planes get coplanar after a flip. See Figure 1.3 for an example. There exist liftings for the pointed pseudo-triangulation on the left side, such that all shown edges appear in the resulting surface. This is not true for the complex at the right that is the result of flipping the bold edge. There exists no lifting such that the dashed edges will show up in the resulting surface. Also, the vertices solely incident to disappearing edges will vanish. Because of this behavior, it can

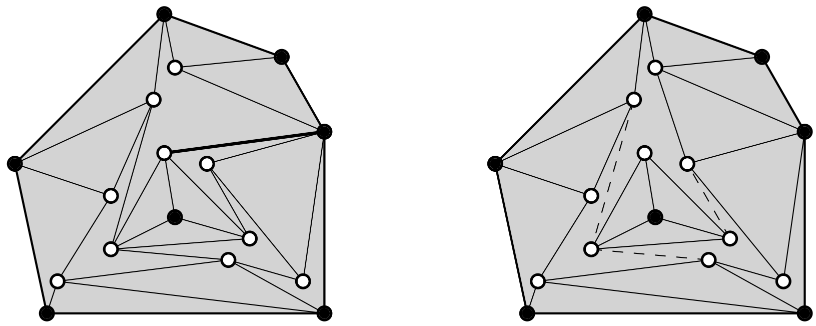

Figure 1.4: Applying a surface flip to an edge (bold) in a pseudo-triangulation (left) can result in a complex (right) that is no pseudo-triangulation any more. The dashed edges do not show up in the surface.

happen that the projection of the resulting surface of a flip is no longer a pseudo-triangulation. Figure 1.4 exemplifies such a situation. A surface flip, applied to the bold edge in the pointed pseudo-triangulation on the left hand side, results in the complex shown on the right hand side.

There exists no lifting such that the dashed edges will appear in the resulting surface. But without these edges the depicted complex is not a pseudo-triangulation. This can be devastating for flip sequences, as valid flips lead to leaving the polygonal complex class where the used flip is no longer defined then. In [52] it is shown that flipping around such situations is always possible, when flipping to locally convex functions. But the freedom of arbitrarily flipping valid edges is lost.

In Chapter 4 we will address this problem. We will introduce so called pre-triangles, which loosely speaking are pseudo-triangles with convex holes. These new faces lead to pre-triangulations, which turn out to be an important structure for flipping in surfaces. A more general setting, namely surfaces on polygonal regions, and a more detailed investigation will lead to further categories of polygonal complexes and ultimately to a new flip operation. This flip is compatible to the standard flips on triangulations and pseudo-triangulations, in the sense that it can simulate their behavior on these structures. But this new flip can also be applied to more general structures. Similar to flips on pseudo-triangulations, the new flip is not a local operation. Even worse, this flip is also not of constant size, as there may be a huge number of edges (at most linear in the number of points, of course) involved. But in some sense, the flip is still confined to a well defined (sub)region.

Parts of this thesis have already been presented at conferences [6, 8, 13] and published in journals [5, 7]. Thus, the content of this work successfully passed already many refereeing stages. These publications are marked bold in the publication list of the appended curriculum vitae (page 65).

1.2 Basic definitions and general notations

Even well known notations have slightly varying definitions and meanings in different areas of research, and sometimes even differ from publication to publication. Therefore we (re)state some of the basic definitions needed in the next chapters. For the last chapter we will, however, need to additionally calibrate some of the notations, as we will have a much more general setting there and thus will need more specific definitions.

In this thesis we will only consider point sets and geometric graphs in the two and three-dimensional Euclidean space, denoted by $\mathbb{E}^2$ and $\mathbb{E}^3$, respectively. Throughout this thesis we will consider point sets, S, in $\mathbb{E}^2$ to be in general position. That is, no three points of S are collinear. In Chapter 4 we will work on point sets in $\mathbb{E}^3$. These three-dimensional point sets are special, as they stem from liftings, where a height vector is applied to a point set in $\mathbb{E}^2$. To put this in other words: the surface resulting from that liftings is always xy-monotone. In Chapter 4 the general position assumption is extended to the predefined height vector. This means that the height vector is given in a way that, for the point set in $\mathbb{E}^2$ (in general position), no four points are defined to be in the same plane in $\mathbb{E}^3$. But, as we will see, a proper lifting is not always applicable with such a height vector, as some vertices are forced into a common plane. And as a last convention on point sets: In many publications, thus also in this work, the cardinality of S is denoted by n.

A *simple polygon* is a subset P of $\mathbb{E}^2$ that is homeomorphic to a disk and has a piecewise linear boundary. The boundary ∂P of P is a connected component of straight line segments that have their endpoints in S. The segments are the *edges* of the polygon and the endpoints of the edges are its *vertices*. As we will consider exclusively simple polygons, we will sometimes refrain from using the term 'simple' for brevity. A simple polygon consists of an open set called *interior* of P, and of ∂P, and partitions $\mathbb{E}^2$ into two sets, a closed set being the polygon itself and an open set called the *exterior* of the polygon. The *size* of a polygon P is the number of its vertices, that is, the number of points of S in ∂P. Obviously, the number of edges of a polygon equals the number of its vertices. Thus, the size of a polygon also refers to the number of its edges. A *diagonal* of a polygon P is a straight line segment which is completely contained in P but is not an edge of P, and whose endpoints are vertices of P. A simple polygon is *convex* if for each vertex of the polygon the angle between the two incident edges is larger than π in the exterior. An equivalent definition for convex polygons requires for each line segment spanned by two vertices to be completely contained in the polygon.

The *convex hull* of S, $CH(S)$, is the smallest convex polygon containing all points of S. The size (or cardinality) of the convex hull is denoted with h. A point set S is said to be in *convex position* if n equals h.

A *triangulation* $T(S)$ of a point set S is a partition of $CH(S)$ into triangles, such that the vertices of $T(S)$ are the points of S. As adding to $T(S)$ any straight line segment between points of S leads to a crossing, triangulations are also referred to as maximal plane straight line graphs. The number of edges and faces of a triangulation depends only on the size of the point set and the cardinality of the convex hull. Using Euler's formula for plane graphs, we easily get $3n-h-3$ for the number of edges and $2n-h-2$ for the number of faces for $T(S)$.

Sometimes notations are slightly adapted to meet the needs for brevity and/or clarity. For instance, if it is clear from the context, a triangulation of a point set can also be denoted by T, without the additional information in the parentheses. On the other hand, the triangulation of a simple polygon P can be denoted by $T(P)$. Such a triangulation uses only the vertices of P as its vertices, and the edges of $T(P)$ have to be completely contained in P. Every simple polygon can be triangulated with $|P|-3$ diagonals. Together with the edges of P, this sums to $2|P|-3$ edges for $T(P)$.

A vertex of a polygon is called a *corner*, if its angle in the interior of the polygon is smaller than π. If the interior angle of a vertex is larger than π, the vertex is denoted a *noncorner*. Remember that we exclude collinear points. Thus, a vertex of a polygon is either a corner or a noncorner of this polygon. In Chapter 4 we will use a slightly stricter definition for corners and noncorners. This will be necessary, as in contrast to simple polygons, a vertex of a polygonal region (which will be defined precisely in Chapter 4) may have many interior and exterior angles.

Clearly, a triangle is a polygon having exactly three vertices, all being corners. A polygon that has exactly three corners, but may have an arbitrary number (including zero) of additional noncorners, is called a *pseudo-triangle*. Pseudo-triangles share many important properties of triangles, while giving the freedom of arbitrary size. Being a simple polygon, a pseudo-triangle ∇ can of course be triangulated. The number of diagonals used for a triangulation equals the number of noncorners of ∇.

Allowing pseudo-triangles makes path for a relaxation of triangulations. Similar to their definition above, a so called *pseudo-triangulation* $PT(S)$ of a point set S is a partition of $CH(S)$ into pseudo-triangles, such that the vertices of $PT(S)$ are the points of S. See [79] for a survey on pseudo-triangulations. Consider a vertex v of $PT(S)$. If each angle between two consecutive edges incident to v is smaller than π, then we name v *non-pointed*. Otherwise, v is called *pointed*. Note that each vertex of $CH(S)$ is pointed for every pseudo-triangulation of S. If all the vertices of a pseudo-triangulation $PT(S)$ that are not vertices of $CH(S)$ are non-pointed, then $PT(S)$ is even a triangulation of S. In this case, the number of non-pointed vertices (and also the number of edges) is maximal. The other extremum, where every vertex of a pseudo-triangulation is pointed, is called a *pointed pseudo-triangulation*. In contrast to triangulations, pointed pseudo-triangulations are the class of pseudo-triangulations that minimize the number of edges for a given point set. Hence, pointed pseudo-triangulations are sometimes also called minimum pseudo-triangulations. We refrain from using this notation, as it may cause confusion in context with minimum weight (pseudo-)triangulations, which we will deal with in the next chapter.

In this section we established some general notation and provided some basic definitions including pseudo-triangulations. In the course of this thesis, it will happen that some of these definitions are given repeatedly (for the convenience of the reader and for the sake of completeness) or even get reformulated in a more precise way. However, we will not redefine previously used definitions this way. We will only give more rigorous (and slightly more complex) definitions to meet the need for more accurate statements.

Chapter 2

Minimum weight pseudo-triangulations

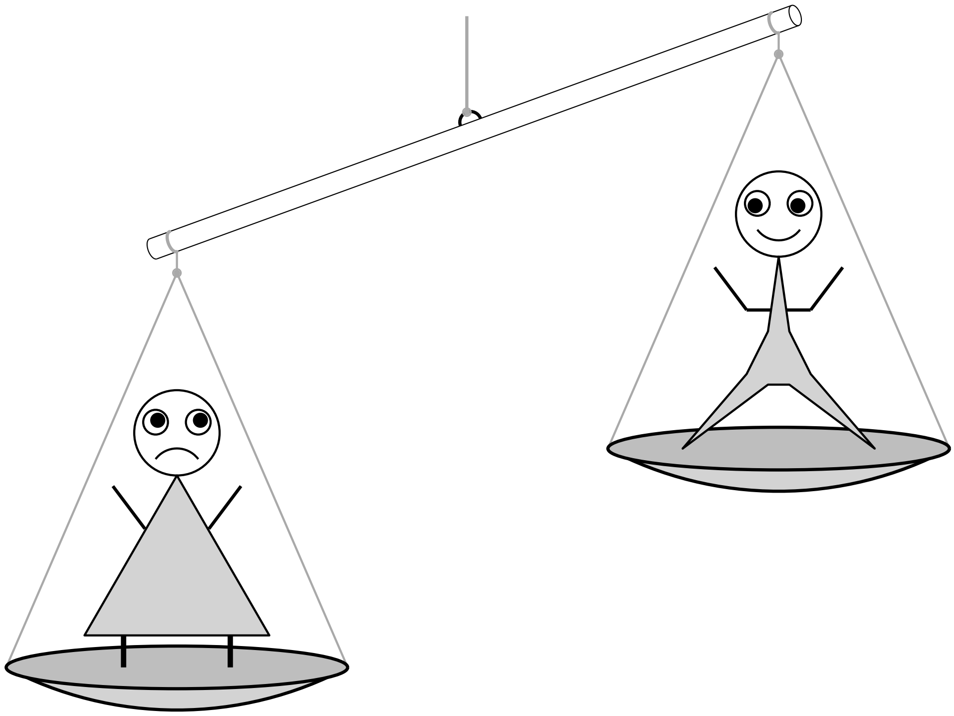

We will begin with some results on optimal pseudo-triangulations, to get the reader started on the concept of pseudo-triangulations. Later, especially in Chapter 4, it will be vitally important to be acquainted with this topic. This work has already been presented at [8] and published in [7].

Optimal triangulations for a set of points in the plane have been, and still are, extensively studied within Computational Geometry. There are many possible optimality criteria, often based on edge lengths or angles. One of the most prominent criteria is the *weight* of a triangulation, that is, the total Euclidean edge length. Computing a *minimum weight triangulation* (MWT) for a point set has been a challenging open problem for many years [45] and various approximation algorithms were proposed over time; see e.g. [26] for a short survey. Mulzer and Rote [70] showed only very recently that the MWT problem is NP-hard.

Also for pseudo-triangulations several optimality criteria have been studied, for example, concerning the maximum face or vertex degree [57]. Optimal pseudo-triangulations can also be found via certain polytope representations [71, 78] or via a realization as locally convex surfaces in three-space [10]. Not all of these optimality criteria have natural counterparts for triangulations. Here we consider the classic minimum weight criterion for pseudo-triangulations.

Rote et al. [81] were the first to ask for an algorithm to compute a minimum weight pseudo-triangulation (MWPT). The complexity of the MWPT problem is unknown, but Levcopoulos and Gudmundsson [50] show that a 12-approximation of an MWPT can be computed in $O(n^3)$ time. Moreover, they give an $O(\log n \cdot w(\text{MST}))$ approximation of an MWPT, in $O(n \log n)$ time. Here $w(\text{MST})$ is the weight of the minimum Euclidean spanning tree, which is a subset of the obtained structure.

Recall from Section 1.2 that a pseudo-triangulation is called *pointed* if every vertex p has one incident region (either a pseudo-triangle or the exterior face) whose angle at p is greater than π. A pointed pseudo-triangulation minimizes the number of edges among all pseudo-triangulations of a given point set. Since a spanning tree of a triangulation is not necessarily pointed (see [15]), the pseudo-triangulation constructed by the approximation algorithm of [50] is also not necessarily pointed. It suggests itself to conjecture that the MWPT should be pointed. However, we show that this does not need to be the case. As a consequence, the MWPT and the minimum weight pointed pseudo-triangulation (MWPPT) of a point set are two different concepts. We also discuss the relation of MWP(P)Ts to greedy pseudo-triangulations, and we give conditions on point sets under which the MWPT is lighter than the MWT.

2.1 Does an MWPT have to be pointed?

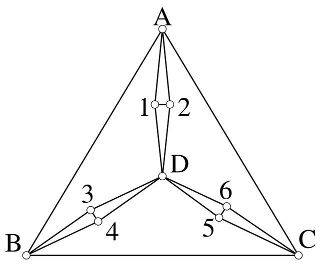

Figure 2.1: Non-pointed MWPT of a point set.

Lemma 1. *The minimum weight pseudo-triangulation of a point set is not necessarily pointed.*

Proof. Consider the pseudo-triangulation, PT, in Figure 2.1, containing exactly one non-pointed vertex. To make PT pointed means to reduce the number of its edges by exactly one. However, PT does not allow to simply remove one single edge to create a pointed pseudo-triangulation. Observe further that all used inner edges adjacent to A, B, C, or D have equal length, and that the edges $\overline{12}$, $\overline{34}$, and $\overline{56}$ are arbitrarily short. It is easy to see that $|\overline{13}| \approx \sqrt{3} \cdot |\overline{1D}|$. Thus the shortest non-used edge $\overline{13}$ (or equivalent) is longer than 3/2 times the length of the longest used edge (apart from convex hull edges, of course). This means, substituting k used edges with $k-1$ unused edges leads to a weight gain for $k > 2$.

It remains to show that no two edges of PT can be replaced by a single new edge without increasing the weight of the pseudo-triangulation. We will exclude the unused edges one by one. Inserting the edge $\overline{AD}$ we have to delete $\overline{12}$ from PT because they would cross. As $\overline{12}$ is arbitrarily short, $\overline{AD}$ is longer, than one longest used edge of PT and $\overline{12}$ together. Edge $\overline{A3}$ is longer than the edges $\overline{A1}$ and $\overline{1D}$ together, which are the longest interior edges of PT. Therefore, $\overline{A3}$ may not be inserted instead of two used interior edges because this would raise the weight of PT. Edge $\overline{A4}$ is inapplicable because it is even longer than edge $\overline{A3}$. Finally, the insertion of edge $\overline{13}$, $\overline{14}$, or $\overline{24}$ would require the insertion of at least one other unused edge, which inevitably leads to a weight gain. The remaining unused edges are equivalent to already mentioned ones. $\square$

2.2 Vertex degrees in an MW(P)PT

For many problems, it would be a major step towards efficient algorithms if a data structure could be found, whose representation as a graph has a low (constant) degree for each of its vertices. For triangulations, there exist simple examples for point sets, where every triangulation has at least one vertex of high (linear) degree. On the other hand, it is known that for every point set there exists a pseudo-triangulation with maximum vertex degree five [57]. Considering optimal structures, the question arises immediately whether pseudo-triangulations convey this property to the minimum weight problem. That is, whether there exists a constant (or at least a sublinear) degree bound for minimum weight (pointed) pseudo-triangulations.

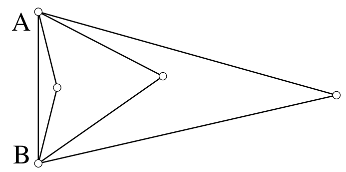

Figure 2.2: An MWPPT with linear vertex degree.

Lemma 2. *A minimum weight (pointed) pseudo-triangulation can have vertices with arbitrarily high vertex degree.*

Proof. See Figure 2.2. For each two consecutive triangles based on $\overline{AB}$, the distance between their tips is larger than the longest edge of the smaller triangle. This implies that the shown pseudo-triangulation is indeed minimum weight. The degree of the vertices A and B is $n-1$ if the example is drawn on n points. □

2.3 Greedy (pointed) pseudo-triangulations

The greedy pseudo-triangulation of a point set S is obtained by inserting edges spanned by S in increasing length order, such that no crossings are caused and until a pseudo-triangulation of S is obtained. Though such a greedy pseudo-triangulation clearly exists, the concept is not meaningful, as we are going to show below.

Lemma 3. *Let ∇ be any pseudo-triangle that is not a triangle. Then ∇ contains some diagonal that is shorter than the longest edge of ∇.*

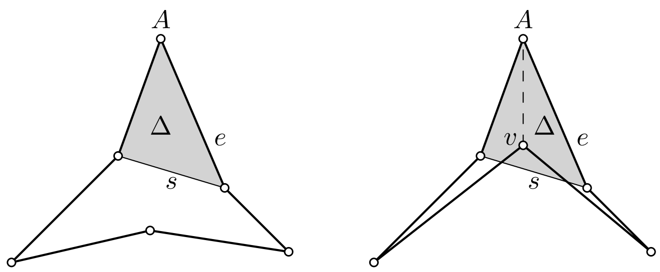

Figure 2.3: Edge lengths inside a pseudo-triangle.

Proof. As the sum of angles in a triangle is π, it is immediate that the three (interior) angles at the corners of ∇ sum up to less than π. Hence there exists a corner A of ∇ where the interior angle is less than $\frac{\pi}{3}$. Let s be the line segment connecting the two vertices of ∇ neighbored to A. Moreover, denote with Δ the triangle spanned by s and A. Clearly, the longest edge of Δ is not s

but rather an edge of ∇, say e. So, if s is a diagonal of ∇ then we are done. See Figure 2.3(left) for an example. Otherwise, there have to exist vertices in the interior of Δ. Corner A sees at least one of them, v, and $\overline{Av}$ is a diagonal of ∇ that is shorter than e. See Figure 2.3(right). The lemma follows. □

Corollary 1. *For every point set S, the greedy pseudo-triangulation equals the greedy triangulation.*

Proof. Assume that the greedy pseudo-triangulation of S contains a pseudo-triangle, ∇, that is not a triangle. Then, by Lemma 3, ∇ contains some diagonal, d, being shorter than its longest edge. So, during the greedy process of constructing the pseudo-triangulation, d would have been inserted before completing the insertion of the edges that form ∇ – a contradiction. □

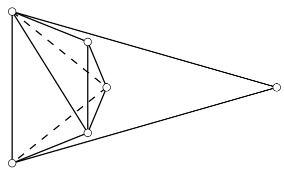

Figure 2.4: The greedy pointed pseudo-triangulation (solid) differs from the MWPPT (dashed).

Requiring pointedness of a greedy pseudo-triangulation changes the situation. This concept is well defined, too, as each face of the pointed graph produced so far – if not a pseudo-triangle – can be split into two faces using any geodesic and without violating pointedness. Not surprisingly, the greedy pointed pseudo-triangulation can differ from the MWPPT. Figure 2.4 gives a simple example. There are sets of n points for which the greedy triangulation has length $\Omega(\sqrt{n})$ times the length of the minimum weight triangulation [68]. This bound is tight [69]. Unfortunately it seems that neither of these constructions is applicable to pseudo-triangulations, which raises the question of how well the greedy pointed pseudo-triangulation approximates the MWPPT.

2.4 Comparing MWPT and MWT

In this section, we compare the minimum weight pseudo-triangulation to the minimum weight triangulation of a point set. The so-called *wheel* will turn out to be a useful structure for this comparison. A wheel is the star-like triangulation of a convex polygon with exactly one interior vertex, the *hub* of the wheel. We call the vertex degree of the hub (that is, the size of the convex polygon) the *degree* of a wheel. The *spokes* of a wheel are the edges of the wheel incident to the hub. Let us call a *big angle* an angle that is larger than π.

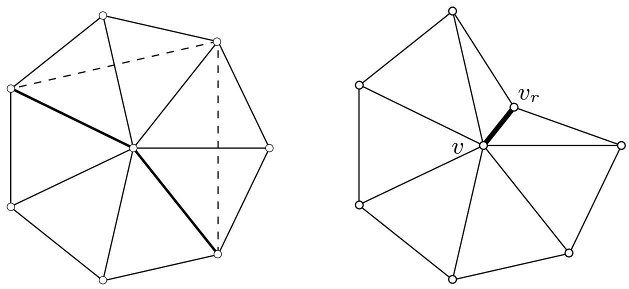

Figure 2.5: Left: This regular wheel of degree 7 is both the MWT and the MWPT of the underlying point set. The dashed segments indicate the two edges that would be needed to make the hub pointed. Right: The bold edge can be removed, as the triangles incident to v do not form a wheel.

Lemma 4. *There are sets of $n \geq 8$ points, that do not lie in convex position, for which the minimum weight pseudo-triangulation is a triangulation.*

Proof. Consider the regular wheel in Figure 2.5(left). It is easy to see that this wheel is the MWT of the underlying point set. To construct a pseudo-triangulation that is not a triangulation we have to make the hub pointed, as it is the only interior vertex. This involves removing at least 3 spokes, and inserting 2 non-spoke edges (dashed edges in Figure 2.5(left)). Let δ be the length of the shortest non-spoke edge and let R be the length of a spoke. We have $\delta = 2\sin(2\pi/7)\ R \approx 1.564\ R$ and hence, any 2 non-spoke edges are longer than 3 times R. $\square$

The point set used in the proof of Lemma 4 contains only one vertex in the interior. Therefore, the question arises whether requiring a certain number of interior vertices in a point set always ensures the existence of pointed vertices in its MWPT. We settle this question in the affirmative in the remainder of this section.

Consider an interior vertex v of a triangulation T. The triangles incident to v form a triangulation $T_v \subseteq T$. The edges on the boundary of T (edges not adjacent to v) form a polygon ∂T_v. If ∂T_v is convex, then T_v is a wheel. Otherwise, there exists at least one vertex v_r which is reflex on ∂T_v. It is easy to see, that the edge connecting v and v_r can be removed, resulting in a pseudo-triangle composed of two triangles, and a new pointed vertex v_r. See Figure 2.5(right) for an example.

Observation 1. *If the MWPT of a point set is a triangulation then, for each interior vertex, its incident triangles form a wheel.*

A polygon in the plane always has at least 3 corners. Recall also that a pseudo-triangle has exactly 3 corners, and that in a pseudo-triangulation internal edges are incident to exactly two faces. Thus, removing an interior edge of a pseudo-triangulation either results in another pseudo-triangulation, or creates a pseudo-quadrilateral (a polygon with exactly 4 corners). The latter happens if and only if the number of pointed vertices does not increase.

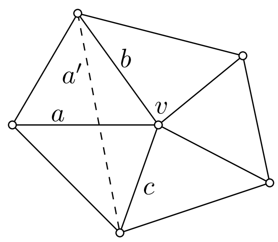

Figure 2.6: A non-regular wheel of degree 5.

Observation 2. *An MWPT cannot have an edge whose removal changes a vertex from non-pointed to pointed.*

It is easy to see that every non-pointed vertex of degree 3 or 4 has at least one adjacent edge whose removal makes this vertex pointed. We can always draw a line through such a vertex such that one of the two half-planes delimited by this line contains only one single edge. This edge can then be removed to make the vertex pointed. Note that this is not necessarily true for non-pointed vertices of degree 5 or higher (for example, consider Figure 2.6). The smallest number of edges that can be guaranteed to lie in one half-plane are $\left\lfloor \frac{k-1}{2} \right\rfloor$, with k being the degree of the vertex. Removing these edges will leave a hole that has to be pseudo-triangulated by inserting the same number of edges minus one. See also (the proofs of) the following Lemmas 5 and 6.

Observation 3. *An MWPT cannot have non-pointed interior vertices of degree less than 5.*

We continue with a series of properties of wheels that imply the property MWPT $\neq$ MWT.

Observation 4. *Let a, b, and c be three spokes of a wheel and let a' be the edge joining the non-hub endpoints of b and c. If $|a| > |c|$ then $|a| + |b| > |a'|$.*

This obviously holds since $|a| + |b| > |c| + |b| > |a'|$. In the following we use Observation 4 only in configurations where a crosses a'.

Lemma 5. *If the minimum weight triangulation of a point set is a wheel of degree* 5 *then there exists a pseudo-triangulation on these 6 points that is lighter.*

Proof. See Figure 2.6 for an example of such a wheel. If there is a spoke whose removal makes the interior vertex v pointed, then Observation 2 implies the lemma. Otherwise, the removal of any pair of consecutive spokes makes v pointed. Let a be the longest spoke and let b and c be its neighbors. Then according to Observation 4 removing a and b and adding a' results in a lighter pseudo-triangulation. $\square$

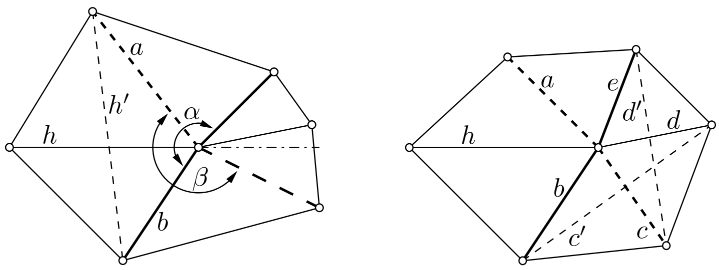

Figure 2.7: Two non-regular wheels of degree 6.

Lemma 6. *If the minimum weight triangulation of a point set is a wheel of degree* 6 *then there exists a pseudo-triangulation on these* 7 *points that is lighter.*

Proof. Let h be the longest spoke. If there is a spoke whose removal makes the interior vertex pointed, then Observation 2 implies the lemma. Otherwise, assuming general position, there are at least 2 spokes on either side of the line supporting h. If $\alpha > \pi$ or $\beta > \pi$ in Figure 2.7 (left) then Observation 4 implies the lemma. So assume $\alpha < \pi$ and $\beta < \pi$ (Figure 2.7 (right)). This implies that removing c and d will make the interior vertex pointed. If $|c| > |d|$ then we add c', otherwise we add d'. Now Observation 4 again implies the lemma. $\square$

Let us summarize: An MWPT can contain a wheel of degree 7 or higher, but not of degree 6 or lower. This directly implies (together with Observation 1) that an MWPT, which is a triangulation, cannot have interior vertices of degree 6 or lower. Further, no MWPT can have non-pointed interior vertices of degree 4 or 3. Let $S_i \subset S$ be the subset of points that are in the interior of the convex hull $CH(S)$ of a point set S. For brevity we will call points in S_i *interior points* (and *interior vertices* for graphs). We can use the summarized facts to give a lower bound on the size of S_i, such that the minimum weight pseudo-triangulation is not a triangulation. To this end we have to give a lower bound for the total degree used on $CH(S)$.

Lemma 7. *Let S be a point set with h points on the convex hull and at least* 3 *interior points. For every triangulation of S the sum of degrees of the convex hull vertices, ρ_h, is at least* $3 \cdot h + 3$.

Proof. It is easy to see that ρ_h is minimized if the interior vertices have a triangular convex hull. Thus we only have to consider h points in convex position and a triangle, Δ, inside. The number of triangulation edges then is $2 \cdot h + 6$, exactly $h + 3$ of which are interior to the belt $conv(S) \setminus \Delta$. But each of these interior edges has to be incident to some vertex of $conv(S)$. We conclude $\rho_h \geq (h + 3) + 2 \cdot h = 3 \cdot h + 3$. $\square$

Theorem 1. *If a set S of $n \geq 15$ points contains more than $\frac{n-9}{2}$ interior points then its minimum weight pseudo-triangulation contains pointed interior vertices.*

Proof. Any triangulation with average *interior* vertex degree $\overline{\rho}_i < 7$ has at least one interior vertex of degree at most 6. From Observations 1 and 3 and Lemmas 5 and 6 we know that, in such a case, we can construct a corresponding pseudo-triangulation which is lighter than this triangulation.

The sum of all vertex degrees in a triangulation of S is exactly $6 \cdot n - 2 \cdot h - 6$ if h points of S are extreme. By Lemma 7, the sum of interior vertex degrees is at most $6 \cdot n - 5 \cdot h - 9$, which gives $n + 5 \cdot i - 9$ if there are $i = n - h \geq 3$ interior points. The average interior vertex degree thus is $\overline{\rho}_i \leq \frac{n-9}{i} + 5$. If we want $\overline{\rho}_i < 7$ then $i > \frac{n-9}{2}$. □

2.5 Chapter summary

We have given some properties of minimum weight pseudo-triangulations. A main open question is how much weight loss can be guaranteed for any set with sufficiently many interior points when relaxing from triangulations to pseudo-triangulations. Lemma 4 shows that there might be no gain at all and, even worse, the MWPPT may be longer than the minimum weight triangulation. On the other hand, Theorem 1 suggests that the gain might be linear in the number of interior vertices.

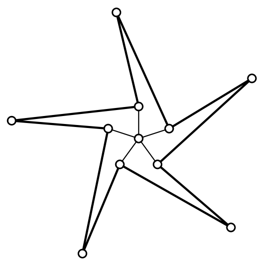

Figure 2.8: Non-pointed MWPT of a pointgon.

In this chapter we considered MWP(P)Ts of point sets. Gudmundsson and Levcopoulos [50] showed how to compute the MWPT or the MWPPT of a simple polygon in cubic time. However, little is known about MWP(P)Ts of so-called *pointgons* – point sets inside simple polygons. Figure 2.8 shows that the MWPT of a pointgon is not necessarily pointed; it can have non-pointed interior vertices of degree 5. For point sets, we know that an MWPT cannot have non-pointed interior vertices of degree 4 and 3. Further, there exist point sets which have non-pointed interior vertices of degree 6 in their MWPT (see Figure 2.1). It would be interesting to compare point sets and pointgons in this respect.

Chapter 3

Flip-graphs of degree-bounded (pseudo) triangulations

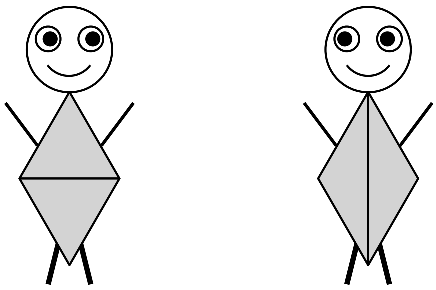

We used the previous chapter to introduce pseudo-triangulations. To consolidate the confidence in handling pseudo-triangulations we compared them to triangulations with respect to the most prominent optimality criterion on these structures, the minimum total edge length. The presented minimum weight (pseudo-)triangulation is but one instance among all the (pseudo-)triangulations. For proving the optimality of particular examples, we actually showed that none of the unused edges could be used to reduce the total edge length. Following the proofs we could actually create the argued structures. This could lead to inserting an edge that crosses many edges of the triangulation. All the intersected edges would have to be removed and the resulting "hole" would have to be triangulated. This can be a major transformation of the initial triangulation, leading to a completely different triangulation that most likely would not represent the single unique result of the operation. Such a transformation is not useful for most applications (if not all).

We utilize the proofs and results of the present chapter to introduce a widely used technique for restructuring graphs. A very general definition for such an operation can be found on the web page ("The flip corner") of Ferran Hurtado [55]: "A flip is a local transformation in which a new object in a class is produced by means of a small modification of a previous object in the same class." Here, "local transformation" means that only a small (preferably constant) part of the object takes part in the flip operation, whereas "small modification" refers to the size of the actual change. Hence, a flip is a practical tool to create different objects of one class or to move from one object towards another special object of the same class with small, carefully defined steps. "It is common to define a graph having a node for each object in the class, adjacencies corresponding to flips. Then connectedness, diameter, shortest paths, and so on, are suitable expressions for natural issues about the flip operation" [55]. Such a graph is often called a flip graph. Moving in this graph on a predefined path can, for example, be used to enumerate (and generate) all different objects of a class.

In this chapter we concentrate on flipping in triangulations, more precise, triangulations on convex point sets with some restrictions on the degree bound. Parts of this work have been presented (as an extended abstract) at [13]. A full version for journal publication is in preparation.

For triangulations, an *edge flip* exchanges a diagonal of a convex quadrilateral, formed by two triangles, with its counterpart. This is a common operation on triangulations, often also referred to as Lawson flip [63]. An edge flip on triangulations is local, as it only affects a convex quadrilateral (which is of constant size) and has constant size itself ("small modification"), as a single edge gets exchanged by another single edge. Furthermore, an edge flip in a triangulation is reversible in the following sense: An edge flip exchanging the edge e with e', directly followed by another edge flip applied to e', results in edge e. See the two flips in Figure 3.1(a). Note that not all edges in a triangulation are flippable. Edges of the convex hull are obviously not flippable, as they are part of every triangulation. The same is true for internal edges that are not crossed by any edge in the complete graph of the underlying point set. Such edges are sometimes called *unavoidable*, see for example [4, 11]. Besides unavoidable edges there exists another class of internal edges, which are not applicable for flipping in triangulations. These are edges whose two incident triangles of the current triangulation do not form a convex quadrilateral; see the bold edge in Figure 3.1(b).

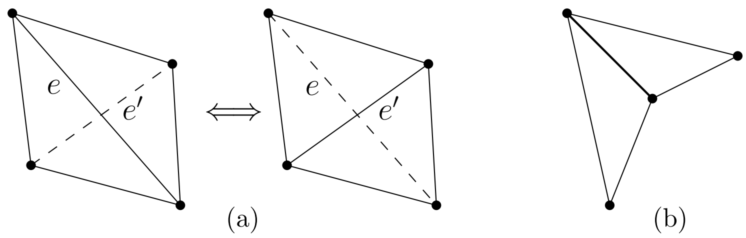

Figure 3.1: (a) The reversibility of the edge flip. Both times, the solid diagonal of the quadrilateral is flipped to the dashed one. (b) An example for a non-flippable edge (bold) in a triangulation.

The *flip graph* $\mathcal{F}_T(S)$ of triangulations of a planar point set S contains a vertex for every triangulation of S, and two vertices are connected by an edge if there is a flip that transforms the corresponding triangulations into each other. One of the first and most fundamental results concerning edge flips in triangulations is the fact that flips can be used repeatedly to convert any triangulation into the Delaunay triangulation [61, 83]. As flips are reversible, this implies immediately that $\mathcal{F}_T(S)$ is connected for any planar point set S.

The *flip distance* between two triangulations is the minimum number of flips needed to convert one triangulation into the other. The diameter of $\mathcal{F}_T(S)$, which is an upper bound on the flip distance, is known to be $\Theta(n)$ if S is in convex position, and $\Theta(n^2)$ if S is in general position. However, computing the flip distance between two particular triangulations has proven to be a difficult problem, and only a few results are known: Hanke et al. [53] showed a lower bound of the flip distance in terms of the number of crossings between the edges of the triangulations, while Eppstein [41] provides an $O(n^2)$ algorithm which computes the flip distance between two triangulations for a class of very special point sets: those having no empty convex pentagon (of course, collinear points are allowed). In general, however, the computational complexity of computing the flip distance between two particular triangulations is not known, even for convex sets [56]. In higher dimensions the flip graph does not even have to be connected [82].

Of growing interest are also subgraphs of flip graphs which correspond to particular classes of triangulations. Houle et al. [54] consider triangulations which contain a perfect matching of the underlying point set. They show that this class of triangulations is connected via flips, that is, the corresponding subgraph of the flip graph is connected. Related results exist for order-k Delaunay graphs, which consist of a subset of k-edges, where a k-edge is an edge for which a covering disk exists which covers at most k other points of the set. For general point sets, the graph of order-k Delaunay graphs is connected via edge flips, for $k \leq 1$, but there exist examples for $k \geq 2$ that can not be converted into each other without leaving this class [1]. If the underlying point set is in convex position, then [1] also shows that the resulting flip graph is connected for any $k \geq 0$. The flip operation has been extended to other planar graphs, see [30] for a very recent and extensive survey.

There are point sets for which every triangulation has a vertex of degree $n-1$, but point sets in convex position always have triangulations with maximum vertex degree 4. Therefore, we concentrate on point sets in convex position and study the following question: is the flip graph of triangulations with maximum vertex degree k connected?

Results. Let S be a set of n points in convex position in the plane. In Section 3.1 we start by exhibiting examples showing that the flip graph of triangulations with maximum vertex degree k cannot be connected if $k \leq 6$. Then we prove that the flip graphs are connected for any $k > 6$ and argue that they have diameter $O(n^2)$. In Section 3.2 we show that we can improve the diameter bound to $O(n \log n)$ if we allow a relaxed degree bound of $k + 4$ in intermediate steps. To this end, we also improve some of the operations used in Section 3.1. Finally, in Section 3.3, we briefly consider point sets in general position and show that flip graphs of triangulations can be disconnected for any constant k.

3.1 Point sets in convex position with tight degree bounds

In this section, we study the flip graphs of bounded degree triangulations of a set S of n points in convex position in the plane. As mentioned above, arbitrary point sets do not necessarily have a triangulation of bounded vertex degree, but point sets in convex position always have triangulations of maximum vertex degree 4. See Figure 3.2(a) for an example. Every point set S in convex position has in fact $\Theta(n)$ different triangulations of this type. It is easy to see, though, that one can not flip even a single edge in such a triangulation without exceeding a vertex degree of 4 .

For maximum degree $k = 5$, consider the triangulation depicted in Figure 3.2(b). Only the dashed edges can be flipped, but there are $\Theta(n)$ rotationally symmetric versions of this triangulation, none of which can be reached from any other without exceeding a vertex degree of 5. For $k = 6$ consider the triangulation depicted in Figure 3.2(c). No edge of this triangulation can be flipped but again there are $\Theta(n)$ rotationally symmetric versions of this triangulation, none of which can be reached from any other without exceeding a vertex degree of 6.

Definitions and notation. Throughout this section and Section 3.2, S will be a set of n points in convex position in the plane. Let D be the dual graph of a triangulation T of S. Because S is in convex position, D is a tree. We distinguish three different types of triangles in T: *ears*, which have two edges on the convex hull of S, *path triangles*, which have one edge on the convex hull of S, and *inner triangles*, which have no edge on the convex hull of S. The *tip* of an ear is the vertex that is adjacent only to two convex hull vertices. The ears of T are dual to the *leaves* of D and inner triangles of T are dual to *branching vertices* of degree three. A *path* in D is any connected subgraph of D that consists only of vertices of degree two or one, so in particular, any vertex of a path is dual to a path triangle or an ear. Note that a vertex of degree one (leaf) will be part of an adjacent path with this definition. A path that contains at least one leaf is called a *leaf path*. All other paths are called *inner paths*. An inner triangle that is adjacent to at least two leaf paths is referred to as *merge triangle*.

We classify the sets of consecutive triangles, that are dual to a path in D, into two possible structures: fans and zigzags. A *fan* is a maximal subset of at least 2 path triangles which is dual to a path, and whose triangles all share one common vertex of S, the *fan handle*. The size of a fan is the number of triangles that make up the fan. Thus, the degree in T of a fan handle is always

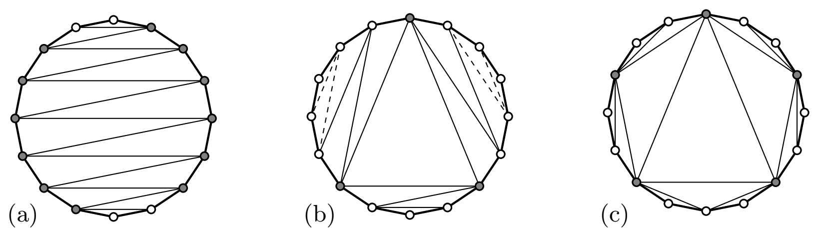

Figure 3.2: Triangulations with vertex degree $k = 4$ (a), $k = 5$ (b), and $k = 6$ (c).

at least five (as the size of a fan is at least 2). The primal of a path in D is said to form a *zigzag* if the set of diagonals is a path. Observe that, in this case, the convex hull edges of every other path triangle are adjacent on the convex hull. Flipping every other diagonal of a zigzag is called an *inversion* of the zigzag.

A triangulation of S is called a *zigzag triangulation* if it has precisely two ears which are connected by a zigzag. Note that the example from Figure 3.2(a) is a zigzag triangulation. Clearly, the maximum vertex degree of a zigzag triangulation is four and its dual graph is a path. A zigzag triangulation is uniquely defined (up to inversion) by the location of one of its ears. We call the zigzag triangulation that has the tip of one of its ears on the left-most vertex of S the *left-most zigzag triangulation* of S.

For brevity, let $\mathcal{T}_k(S)$ be the set of triangulations of S, such that all points have degree at most k. If clear from the context, the reference to S will be omitted. Finally, $T \in \mathcal{T}_k(S)$ is a *fringe triangulation*, if every leaf path is dual to a zigzag. In particular, every zigzag triangulation is a fringe triangulation.

Algorithmic outline. Let $T \in \mathcal{T}_k(S)$ be a triangulation of S with maximum vertex degree $k > 6$. In the following, we show how to flip from T to the left-most zigzag triangulation of S without ever exceeding vertex degree k. In particular, we show in Subsection 3.1.1 how to first convert T into a fringe triangulation by eliminating fans in leaf paths. This does not increase the number of inner triangles. In Subsection 3.1.2 we prove that each fringe triangulation always has a *light* merge triangle, that is, a merge triangle with at least two vertices of degree $< k$. We then show how to remove this light merge triangle by merging its adjacent zigzags with $O(n)$ flips, resulting in a triangulation that has one less inner triangle. This triangulation can again be converted into a fringe triangulation. After repeating this step $O(n)$ times, we have converted T into a zigzag triangulation. The number of flips for eliminating fans to get fringe triangulations is $O(n^2)$ in total. Finally, in Subsection 3.1.3 we demonstrate how to "rotate" any zigzag triangulation into the left-most zigzag triangulation of S using $O(n^2)$ flips.

Theorem 2. *Let S be a set of n points in convex position and let $T \in \mathcal{T}_k(S)$ be a triangulation of S with maximum vertex degree $k > 6$. Then T can be flipped into the left-most zigzag triangulation of S in $O(n^2)$ flips while at no time exceeding a vertex degree of k.*

Corollary 2. *Let S be a set of n points in convex position. Then for any $k > 6$ the set $\mathcal{T}_k(S)$ is connected by flips. The diameter of the corresponding flip graph is $O(n^2)$.*

3.1.1 Creating a fringe triangulation

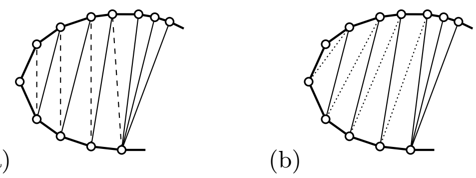

Figure 3.3: A fan adjacent to a zigzag. Inverting the zigzag (flipping the dashed edges to the dotted edges) reduces the size of the fan.

Recall that in a fringe triangulation every leaf path is dual to a zigzag. If T is not a fringe triangulation, then it has at least one fan in a leaf path. Let F be a fan closest to an ear in such a leaf path. Here, closest means that there is no other fan between F and the ear. Let $2 \leq f \leq n-4$ be the size of F and let Z be the zigzag, which is dual to a leaf path and adjacent to F (see Figure 3.3(a)). One inversion of Z can decrease f by one (see Figure 3.3(b)). After the inversion, the reduced fan is again adjacent to a zigzag which is dual to a leaf path. Repeated inversion of the (stepwise expanded) zigzag will eventually remove the fan at a cost of $\Theta(f \cdot (\widetilde{n} + f))$ flips, where $\widetilde{n}$ is the length of the dual path of Z. Observe that every fan will eventually become adjacent to a zigzag, that is dual to a leaf path. Thus the overall cost of removing all fans is $O(n^2)$, as the sum over the size of all fans is $O(n)$.

Lemma 8. *Let S be a set of n points in convex position and let $T \in \mathcal{T}_k(S)$. Then T can be transformed into a fringe triangulation of S in $O(n^2)$ flips, while at no time exceeding a vertex degree of $k > 6$.*

3.1.2 Merging zigzags

Recall that a merge triangle is an inner triangle that is adjacent to at least two leaf paths. In a fringe triangulation, all leaf paths are dual to a zigzag. In the following, we first prove that each fringe triangulation has a light merge triangle, that is, a merge triangle with two vertices of degree $< k$. Then we show how to remove a light merge triangle by merging two of the adjacent zigzags.

Let $\triangle$ be a merge triangle and Z_1, Z_2 be the two incident zigzags dual to a leaf path. We denote by v_{tip} the common vertex of $\triangle$, Z_1, and Z_2 and call v_{tip} the *tip* of $\triangle$. The *base edge* of $\triangle$ is the edge, e, of $\triangle$ that is not adjacent to v_{tip}. See Figure 3.6(a).

Lemma 9. *Let S be a set of n points in convex position and let $T \in \mathcal{T}_k(S)$ be a fringe triangulation of S with maximum vertex degree $k > 6$. Then T has a light merge triangle.*

Proof. Since T is not a zigzag triangulation, it has at least one inner triangle. As the dual graph D of T is a tree, at least one of the inner triangles in T is a merge triangle, $\triangle$, which is adjacent to two zigzags Z_1, Z_2 which are dual to leaf paths. The tip v_{tip} of each merge triangle has degree at most 6: Two edges from $\triangle$, two convex hull edges, and at most one edge each from Z_1 and Z_2. Thus it remains to prove that there exists a merge triangle whose base edge e has a vertex with degree $< k$.

If there is a merge triangle where v_{tip} has degree < 6, then we can invert Z_1 or Z_2 and thereby decrease the degree of at least one endpoint of e. Since all vertices of a zigzag, except v_{tip} and the endpoints of e, have degree at most 4, this inversion maintains the degree bound. Hence we can assume that the tip of each merge triangle has degree 6. Let D be the dual graph of T. We create the graph D' by removing all leaf paths from D. Observe that each triangle $\triangle'$ which is dual to a leaf vertex of D' is a merge triangle in T. We now create a second graph D'' by removing each leaf of D' (see Figure 3.4).

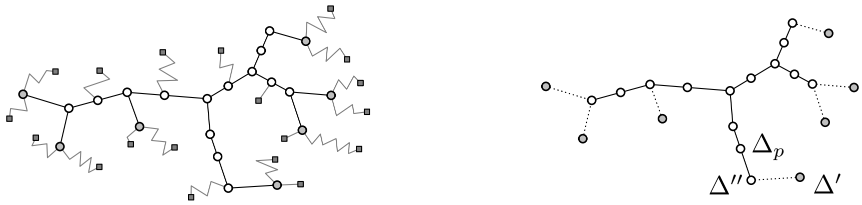

(a) (b)

Figure 3.4: (a) In graph D internal paths are shown as single edges; (b) Graphs D' and D'' (without gray nodes and dotted edges).

Since D is a tree, both D' and D'' are trees as well. Hence D'' has at least two leaves. Let $\triangle''$ be dual to a leaf of D''. Name the only triangle, $\triangle_p$, adjacent to $\triangle''$ in D'' the parent of $\triangle''$. Likewise name the (at most 2) triangle(s), different from $\triangle_p$ and adjacent to $\triangle''$ in D', the children of $\triangle''$.

Let $\triangle'$ (dual to a leaf of D') be a child of $\triangle''$ in D'. Consider the vertex v which is a vertex of both $\triangle''$ and $\triangle'$, but not a vertex of $\triangle_p$ (see Figure 3.5(a) and (b)). By assumption, the tip w of $\triangle'$ has degree 6, which implies that the zigzag adjacent to the edge vw ends at w, and thus no diagonal of the zigzag is adjacent to v. Hence, if $\triangle''$ has only one child in D' (see Figure 3.5(a)) then v has at most degree 6. If $\triangle''$ has two children in D' (see Figure 3.5(b)) then also consider the tip u of $\triangle'$'s sibling. Since u also has degree 6, we again know that the zigzag adjacent to the edge uv ends at u. So, also in this case, we can conclude that v has at most degree 6. $\square$

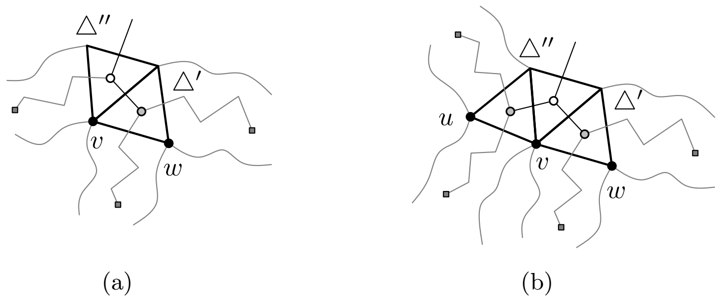

Figure 3.5: The triangle, $\triangle''$, dual to a leaf of D'' can have 1 (a) or 2 (b) children in D'.

Lemma 9 states that each fringe triangulation has a light merge triangle $\triangle$. We now show how to remove $\triangle$ by merging its adjacent zigzags Z_1 and Z_2. Let us denote the edge of $\triangle$ that is adjacent to Z_1 with e_1 and the one adjacent to Z_2 with e_2. Further, let us denote the vertex of $\triangle$ that is shared by e and e_1 by v_1 and the one shared by e and e_2 by v_2 (see Figure 3.6(a)). We can assume w.l.o.g. that v_{tip} has degree 6, that v_1 has degree k, and that v_2 has degree $k-1$.

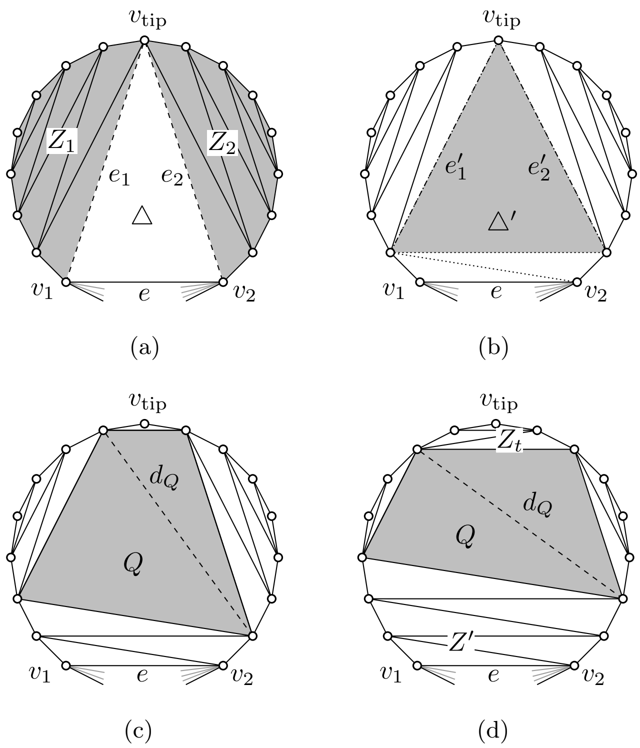

Figure 3.6: Merging two zigzags: Start configuration (a) and the situation after the first (b) and second step (c). Another 2 steps (= 6 flips) lead to the triangulation in (d), giving a good example of the generic configuration.

We want to merge Z_1 and Z_2 together with $\triangle$ into a new zigzag Z' that starts at e. We start by flipping first e_1 and then e_2. These flips create the first two edges of Z' and do not violate the degree bound (see Figure 3.6(b)). Now there is a new triangle $\triangle'$ that lies between (the remains of) Z_1 and Z_2, with edges e'_1 and e'_2, similar to $\triangle$ before. Observe that the two vertices v_1 and v_2, which may have high vertex degree, have been "cut off". For the remaining vertices, we define the following invariant: Each vertex is adjacent to 2 convex hull edges, at most 2 zigzags, and possibly one diagonal, which we will introduce in the next step. If this invariant holds, the degree of each vertex is at most $2 + 2 \cdot 2 + 1 = 7$. Note that this invariant holds for $\triangle'$ as only two of the three zigzags (Z_1, Z_2, Z') can meet at one vertex (of $\triangle'$) and the diagonal has not been introduced yet.

We continue by flipping e'_1. The resulting edge we denote by d_Q. Flipping e'_2 and the edge of Z_1 that was incident to e'_1 concludes this second step. We created a merge quadrilateral Q with the announced diagonal d_Q that separates four zigzags: The shrinking zigzags Z_1 and Z_2, the zigzag Z' growing from e, and a temporary zigzag Z_t growing from v_{tip} (see Figure 3.6(c)).

Flipping d_Q followed by the two edges common to Q, and Z_1 and Z_2, respectively, results in one additional edge for each, Z' and Z_t, while Z_1 and Z_2 get reduced by the same amount. Thus 3 flips per step increase Z' by one triangle. Figure 3.6(d) shows the triangulation of Figure 3.6(c) after 2 such steps. Observe that during all flips a vertex can be incident to at most two zigzags which meet at Q (or a merge triangle). As there only exists at most one diagonal per step, the invariant holds.

If Z_1 and Z_2 have the same size then Z' and Z_t eventually grow into each other, forming one zigzag. This completes the merge. If Z_1 and Z_2 are not of equal size, then the smaller one will disappear after sufficiently many steps. This will cause the quadrilateral Q to collapse to a triangle, $\triangle''$ similar to $\triangle'$, and the diagonal d_Q to disappear. It is easy to see that the invariant holds for this new merge triangle $\triangle''$, which separates three zigzags.

We continue by again introducing the quadrilateral (which also increases Z' by one edge), followed by steps of flipping its diagonal and two zigzag edges, which again increases Z' by one edge per step.

As each step consists of at most 3 flips, one of which adding an edge to our target zigzag Z', the process ends after a linear number of flips. And as the invariant holds in each step, the vertex degree never rises above 7.

We have just described an algorithm for merging 2 zigzags at a merge triangle, which proves the following lemma:

Lemma 10. *Let S be a set of n points in convex position, let $T \in \mathcal{T}_k(S)$ be a fringe triangulation of S, let $\triangle$ be a light merge triangle of T, and let $\widetilde{n}$ be the total number of vertices of $\triangle$ and its two adjacent zigzags that end in ears. Both zigzags and $\triangle$ can be merged into one zigzag ending in an ear with $O(\widetilde{n})$ flips, while at no time exceeding a vertex degree of $k > 6$.*

After removing a light merge triangle $\triangle$ of a fringe triangulation, by merging it with its two adjacent zigzags to a new zigzag Z', there might exist a number of fans between Z' and the next inner triangle. To get again a fringe triangulation, these fans have to be removed. This is done by inverting Z' and can take as much as $O(n^2)$ flips, if the fan is of linear size. On the other hand, the overall size of all fans together is also only linear. And as no fan ever gets increased or newly created (outside the scope of the merge operation) the overall number of flips for removing all fans is also only $O(n^2)$.

3.1.3 Rotating a zigzag triangulation

Recall that a zigzag triangulation is uniquely defined (up to inversion) by the location of one if its ears. A *rotation* is a sequence of flips that transform one zigzag triangulation into another one. In this section we show how to rotate any zigzag triangulation of S into the left-most zigzag triangulation with $O(n^2)$ flips.

Observe that two consecutive inversions of a zigzag (flipping first one half of the edges, then the other half) rotates a zigzag (its ears) by one vertex (see Figure 3.7). As one inversion needs a linear number of flips, we can flip any zigzag triangulation to the left-most zigzag triangulation in $O(n^2)$ flips. Clearly, the degree of a vertex during this rotation is at most 5.

Lemma 11. *Let S be a set of n points in convex position and let T be any zigzag triangulation of S. T can be rotated into the left-most zigzag triangulation of S with $O(n^2)$ flips, while at no time exceeding a vertex degree of $k > 6$.*

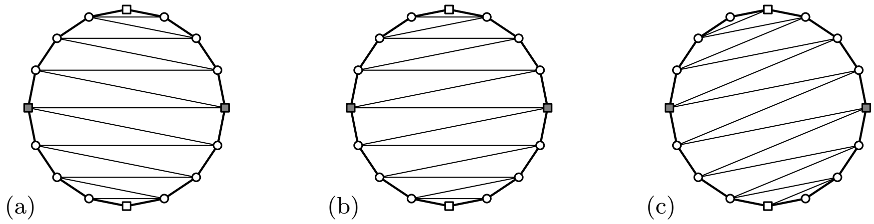

Figure 3.7: Zigzag rotation by one vertex (counter clockwise) using inversion.

3.1.4 Putting things together

With Lemma 11 we are ready to present the proof for Theorem 2.

Proof of Theorem 2. Lemma 8 states that any triangulation with maximum vertex degree $k > 6$ can be transformed into a fringe triangulation with $O(n^2)$ flips. A fringe triangulation obviously contains only a linear number of inner triangles. At least one of them is a light merge triangle, by Lemma 9. Removing light merge triangles repeatedly results in a zigzag triangulation. This takes $O(n)$ flips per triangle, according to Lemma 10, and thus sums up to a total of $O(n^2)$ flips. Lemma 11 concludes the proof by showing that each zigzag triangulation can be rotated to the left-most zigzag triangulation using $O(n^2)$ flips. In addition, each used operation (set of flips) respects the degree bound, as also stated in the corresponding lemmas. □

3.2 Point sets in convex position with relaxed degree bounds

If we allow a slightly increased degree bound in intermediate steps (when flipping from one degree-bounded triangulation to another) then we can lower the diameter of the flip graph to $O(n \log n)$ flips. To this end, we will show that in this relaxed setting we can remove a constant fraction of the inner triangles per step, involving a linear number of flips. For this we need to slightly redefine fringe triangulations, and improve the handling of fans and the rotation of zigzag triangulations. Observe that, if we allow an intermediate degree of at least $k + 1$, then every merge triangle can be merged.

3.2.1 Fringe triangulation redefined

From now on let a *fringe triangulation* T, $T \in \mathcal{T}_k(S)$, be a triangulation which has no fans of size greater than three, where each fan is adjacent to an inner triangle, and where every leaf path is dual to a zigzag. This redefinition puts a limit to the number and the size of fans. For Lemma 8 to remain true for this updated definition, we have to present a more efficient method for removing or reducing fans.

3.2.2 Improved fan handling

For each path we start processing fans at one end, and proceed to the other end. Consider a fan, F, of size f and with handle v^*, like the one depicted in Figure 3.8(a). Let v_f be the first vertex of the fan in the direction of fan processing, and v_0 the last vertex in this direction, respectively.

Because fans are processed in order, all fans preceding F have already been processed, and thus only v_0 can be the fan handle of another fan. In this case, flip the diagonal v^*v_0. This results in both fans being separated by a zigzag and having their size reduced by one. See Figure 3.8(b) for an example where the dashed edge shows the situation before, and the dotted edge after the flip. Thus we can assume that v_0 is not the fan handle of another fan, too.

If neither v_0 nor v_f are vertices of an inner triangle, then their degree is at most 4. Therefore we can convert F to an inner triangle $\triangle(v^*v_fv_0)$ and a zigzag (which might be empty) ending in an ear, performing $O(f)$ flips (see Figure 3.8(a)).

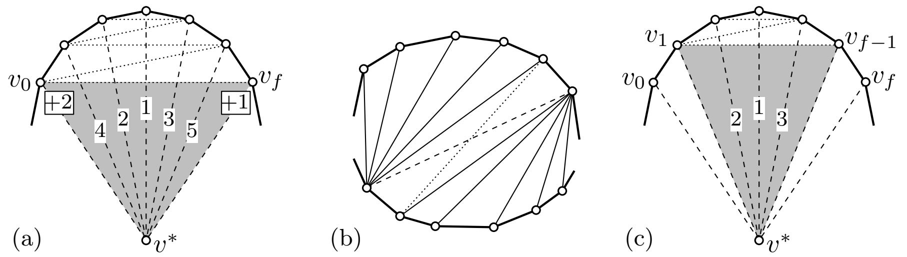

Figure 3.8: Handling fans: The boxed numbers show the degree change, the other numbers indicate the order in which flips are performed.

Let v_0 (or v_f) be a vertex of an inner triangle. If one of v_1 and v_f has degree less than k and the other one less than $k-1$, then we can convert the fan to an inner triangle $\triangle(v^* v_f v_0)$ like above. Suppose this degree requirement is not fulfilled. If $f > 3$ we can convert the fan to an inner triangle $\triangle(v^* v_{f-1} v_1)$ and a zigzag, as both, v_{f-1} and v_1, have degree 3, see Figure 3.8(c). Otherwise, we can keep F, because a fan of size $f \leq 3$ that is adjacent to an inner triangle is allowed in a fringe triangulation.

At last, if all paths have been processed and a constant sized fan remained in a leaf path, then it can be removed by a number of flips linear in the size of the path (a constant number of inversions).

It is easy to see that constructing a redefined fringe triangulation still only takes $O(n)$ flips. Thus Lemma 8 holds. New (additional) inner triangles get created now, but only in the pre-processing step of the initial creation of the fringe triangulation. Unlike before, after a merge operation we now only run into at most two constant size fans. These two fans can be removed with a constant number of inversions, resulting in a number of flips linear in the size of the leaf path. These flips can be accounted to the next merge operation, in which the resulting zigzag takes part.

Observe that with the presented improved fan handling, Lemma 8 not only remains true for the new definition of fringe triangulation, but even improves to the following lemma:

Lemma 12. *Let S be a set of n points in convex position and let $T \in \mathcal{T}_k(S)$. Then T can be transformed into a fringe triangulation of S in $O(n)$ flips, while at no time exceeding a vertex degree of $k > 6$.*

3.2.3 Improved rotation of zigzags

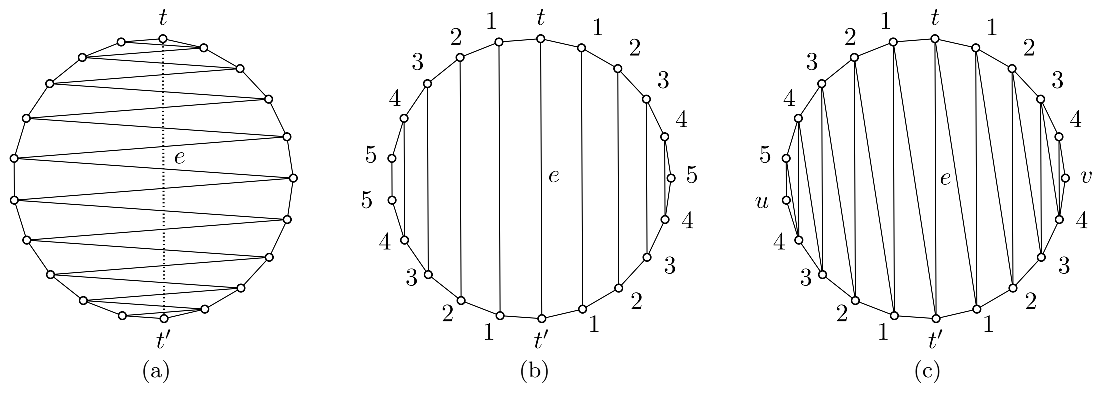

Figure 3.9: Definition of a zigzag triangulation, T^n, normal to the zigzag triangulation, T. (a) initial zigzag triangulation T, with tips of ears t and t' and their connection e; (b) vertex labeling and the edges between equally labeled vertices; (c) the complete zigzag triangulation T^n with tips of ears u and v.

In the next step, merging two inner triangles to one, we will need a rotation operation that needs only a linear number of flips. Of course, such an improved operation will also come in handy for the last step of rotating a zigzag triangulation to the left-most zigzag triangulation. So our next goal is to prove the following statement:

Lemma 13. *Let S be a set of n points in convex position and let T be any zigzag triangulation of S. T can be rotated into any other zigzag triangulation of S with $O(n)$ flips, while at no time exceeding a vertex degree of $k > 6$.*

Proof. Let t and t' be the ears of the initial zigzag triangulation T. First we will look at a special case. We will construct a zigzag triangulation, T^n, that we will call to be *normal* to T. W.l.o.g. let the edge $e = tt'$ be vertical. See Figure 3.9(a). On the left side of e, label the vertices with the minimum distance (along the convex hull) to t or t'. Then connect vertices with equal labels with an edge. Proceed similarly on the right side of e, see Figure 3.9(b). Complete the new zigzag (choose one of the two possibilities) by drawing the remaining diagonals to get T^n. Let the tips of the ears of T^n be u and v, see Figure 3.9(c). Observe that T^n is unique up to an inversion, and that uv is a diagonal in T (up to an inversion).

To flip from T to T^n we first flip the edge connecting u and v to get the edge d_Q. Next flip the other two edges incident to u and v, respectively. This results in a merge quadrilateral Q with diagonal d_Q, splitting the already growing zigzags from u and v from the shrinking zigzags from t and t'. Concerning degree bound and number of flips, the argumentation from Section 3.1.2 applies here as well. Thus, it can be seen that T can be rotated to T^n with $O(n)$ flips, while at no time exceeding a vertex degree of k.

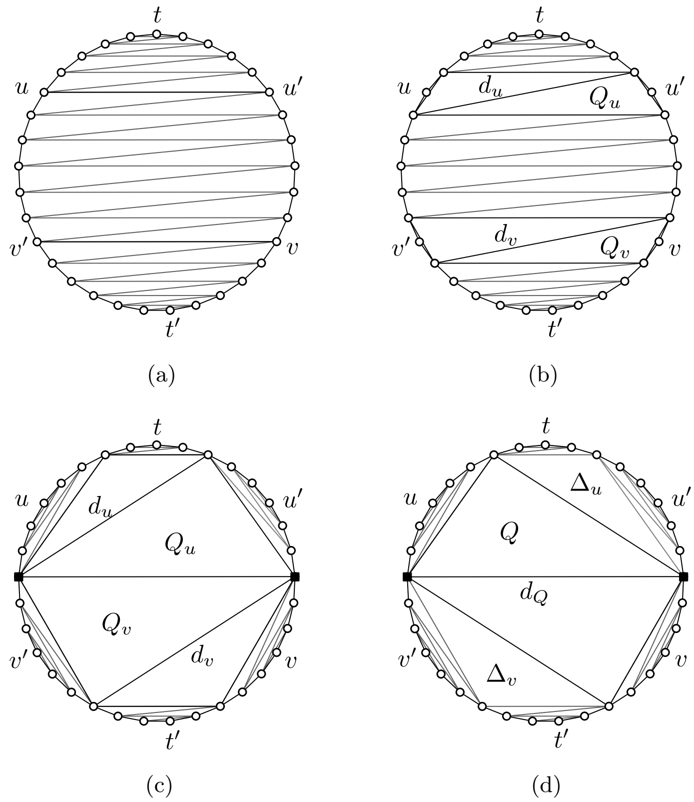

Figure 3.10: General rotation. (a) The initial triangulation T with ears t and t', the ears u and v of the target zigzag triangulation T', and the temporary ears u' and v'. (b) Merge quadrilaterals Q_u and Q_v with diagonals d_u and d_v, respectively. (c) Q_u and Q_v meet and (d) transform into the merge quadrilateral Q with diagonal d_Q and the two merge triangles $\triangle_u$ and $\triangle_v$.

For the general case, the procedure is slightly more involved. Let u and v be the ears of the target zigzag triangulation T'. If T' is not normal to the initial zigzag triangulation T, then u and v are not connected with an edge of T. Let u' (v') be a vertex connected to u (v) with an edge of T, such that the minimum distance along the convex hull between u and v', and between u' and v is equal. See Figure 3.10(a). Flipping the edge connecting u and u' results in the edge d_u. Flipping the other two edges incident to u and u', respectively, results in a merge quadrilateral Q_u, separating the growing zigzags (Z_u and $Z_{u'}$) emanating from u and u' from two shrinking zigzags. The merge quadrilateral Q_v with diagonal d_v (and zigzags Z_v and $Z_{v'}$) is obtained in a similar way. See Figure 3.10(b).

Flipping proceeds with first making a step with Q_u (flipping d_u, followed by a flip to add an edge to Z_u, and another one to add an edge to $Z_{u'}$) followed by a step with Q_v. As long as the two merge quadrilaterals do not meet (share a vertex or an edge), the argumentation from Section 3.1.2 applies, according to the degree bound and also the number of flips.

In the situation when Q_u and Q_v meet, we have to take special care of the two vertices, w and x (marked as black squares in Figure 3.10), which become incident to both quadrilaterals. First observe that the number of steps made with Q_u and Q_v are equal. (In the only situation where this wouldn't be the case, u and v' (and also u' and v) are neighbored on the convex hull. But in this case the zigzag triangulation with ear u is normal to T.) Therefore d_u and d_v have no common vertex when Q_u and Q_v meet (and also not before). For the same reason, only one of Z_u and $Z_{v'}$ adds a degree of 2 to w. Thus w has a degree of at most $2+2+1+1+1 = 7$, with 2 edges from the convex hull, 2 edges from one zigzag and 1 edge from the other, only 1 of d_u and d_v, and 1 remaining edge of T. Similar arguments apply to x. See Figure 3.10(c). Note that in the last step made with Q_v the order of the 3 flips is not arbitrary. To respect the degree bound of w and x, after flipping d_v it may be necessary to first apply the flip adding an edge to Z_v and then to $Z_{v'}$, or vice versa.

Next we invert $Z_{u'}$ and $Z_{v'}$ if they add a degree of 2 to w and x, respectively. Otherwise, they have already the desired direction. Now there either already exists a merge quadrilateral Q, adjacent to Z_u and Z_v (see Figure 3.10(d)), or we have to flip d_u and d_v (like in the example in Figure 3.10(c)). If we had to invert $Z_{u'}$ and $Z_{v'}$ then we can apply both flips. Otherwise, we need to temporarily invert Z_v to reduce the degree of x to 6. Then we can flip d_u (not incident to x before the flip), followed by d_v. Finally, we undo the inversion of Z_v. In total, this is only a constant number of additional inversions.

Introducing the merge quadrilateral Q with diagonal d_Q further results in two light merge triangles $\triangle_u$ and $\triangle_v$, both adjacent to Q and to two zigzags. See Figure 3.10(d). The rest is simple. Merge the two zigzags at $\triangle_u$ and $\triangle_v$, respectively. And then finish with the merge quadrilateral Q like in the case of flipping to T^n.

If Q_u and Q_v do not meet, then they turn into merge triangles because the zigzags from t and t' vanished. This case could be solved similar to the one where Q_u and Q_v meet. Again, in this case there exist vertices that need special care with respect to the degree bound. But unlike before the solution is much more involved. So to simplify the solution in this case we convert it into the already solved case where the two merge quadrilaterals meet. To this end, we rotate the initial zigzag triangulation, T, to the zigzag triangulation, T^n, that is normal to T. Then we rotate T^n to the target zigzag triangulation T', which leads to the desired case. □

3.2.4 Merge two inner triangles to one

Consider the dual graph D of a fringe triangulation, T. We get the reduced graph D' by removing all leaf paths. Each leaf in D' is a merge triangle. Observe that in the relaxed setting all merge triangles are mergeable. Inner triangles, that are dual to nodes of degree 2 in D', are incident to one leaf path in T. Remember that all leaf paths are zigzags. Let two of these inner triangles be *neighbored*, if no other inner triangle exists in the path between them. We will show how to merge two such neighbored inner triangles to one, thereby reducing the total number of inner triangles in T by one.

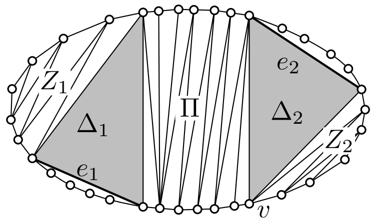

Figure 3.11: Example of two neighbored inner triangles, $\triangle_1$ and $\triangle_2$, with degree 2 in the reduced dual graph D'.

Let $\triangle_1$ and $\triangle_2$ be two neighbored inner triangles that have degree 2 in D'. Let Π be the path between $\triangle_1$ and $\triangle_2$, and let Z_1 (Z_2) be the zigzag dual to the leaf path incident to $\triangle_1$ ($\triangle_2$). Further, let e_1, e_2 be the edges of $\triangle_1$ and $\triangle_2$, respectively, that are not part of a triangle of Π, Z_1, or Z_2. See Figure 3.11 for an example. The diagonals e_1 and e_2 delimit a connected subset $\widetilde{T}$ of T, which is the union of $\triangle_1$, $\triangle_2$, and the triangles of Z_1, Z_2, and Π. We denote the size of $\widetilde{T}$ with $\widetilde{n}$. The following lemma states that we can combine $\triangle_1$ and $\triangle_2$ to (at most one) inner triangle, such that the resulting triangulation is again a fringe triangulation.

Lemma 14. *Given a subset $\widetilde{T}$ of a fringe triangulation T as defined in the above paragraph. We can transform T into a fringe triangulation, such that the two inner triangles $\triangle_1$ and $\triangle_2$ get merged to (at most) one inner triangle. This reduces the number of inner triangles by (at least) 1, takes $O(\widetilde{n})$ flips, while at no time exceeding a vertex degree of $k+4$.*

Proof. We only consider the subset $\widetilde{T}$ which can be obtained by cutting T along e_1 and e_2, turning the diagonals into convex hull edges of $\widetilde{T}$. Observe that $\widetilde{T}$ is dual to a path. As T is a fringe triangulation, $\widetilde{T}$ only contains at most two constant sized fans. Thus, $\widetilde{T}$ can be converted into a zigzag triangulation with a constant number of inversions, at a total of $O(\widetilde{n})$ flips. This zigzag triangulation can be rotated such that it starts at e_1. By Lemma 13, only $O(\widetilde{n})$ flips are needed for this rotation.

This transforms T into a fringe triangulation, T'. Between e_1 and e_2, T' has a zigzag starting at e_1 and ending at an inner triangle, $\triangle$, at e_2, and another zigzag dual to a leaf path incident to $\triangle$. Observe that, if the chains of the convex hull of $\widetilde{T}$ between e_1 and e_2 are of equal size (up to ± 1), then the resulting triangulation between e_1 and e_2 is a zigzag (without inner triangles).

Concerning the degree bound, observe that the used operations, rotation and inversion of a zigzag, always respect the degree bound (inside $\widetilde{T}$). Therefore, the only vertices that can accumulate a degree bigger than k are the vertices of e_1 and e_2. These vertices have a degree of at least 3 inside $\widetilde{T}$ before the transformation. $\widetilde{T}$ contains only fans of constant size (less than 4), and the operations used (inversion, rotation) generate a degree of at most 7 in the present case. Thus, during the transformation, the vertices of e_1 and e_2 get an overall degree in T of at most $k+4$. The result of the merging is again a fringe triangulation with vertex degree at most k, as the result of transforming $\widetilde{T}$ is a zigzag triangulation. (If one of the vertices of e_1 and e_2 had a degree of 3, it may be necessary to invert one (or both) of the zigzags in T' between e_1 and e_2 to restore the vertex degree bound. This only takes another $O(\widetilde{n})$ flips.) □

3.2.5 Putting things together

We now have all tools to show an upper bound of $O(n \log n)$ for the flip distance of bounded degree triangulations with relaxed intermediate degree bounds.

Theorem 3. *Let S be a set of n points in convex position, and let $T \in \mathcal{T}_k(S)$ be a triangulation of S with maximum vertex degree $k > 6$. Then T can be flipped into the left-most zigzag triangulation of S in $O(n \log n)$ flips, while at no time exceeding a vertex degree of $k+4$.*

Proof. By Lemma 12, we can transform T into a fringe triangulation T' as refined in Section 3.2.1. We will show that we always can remove a constant fraction of the inner triangles of T' with a linear number of flips, the result being a fringe triangulation again.

Recall that the dual graph, D, of a (fringe) triangulation of a convex set is always a tree. This is also true for the reduced graph D', where all leaf paths have been cut off in D. All inner triangles dual to degree one nodes (leafs) in D' can be removed. These are merge triangles, and as we allow a degree bound of $k + 4$, they all are light and can be merged to a zigzag, by Lemma 10. It is easy to see that only $O(n)$ flips are needed in total.

Inner triangles dual to degree three nodes in D' can not be removed. But neighbored pairs of inner triangles, dual to nodes of degree 2 in D', can be merged to (at most) one such inner triangle, by Lemma 14.

Let d_1, d_2, d_3 be the number of degree 1, 2, 3 nodes in D', which are the duals of inner triangles in T'. Obviously $d_1 = 2 + d_3$. The number of unmergeable (because they have no neighbor) inner triangles dual to nodes of degree 2 is at most $1 + 2 \cdot d_3$. Therefore, there exists a sufficient number of merge triangles if not enough neighbored inner triangles can be merged, and vice versa. Thus we can reduce the number of inner triangles by a constant fraction with a linear number of flips. This gives an overall bound of $O(n \log n)$ flips. □

3.3 Point sets in general position

In this section we study flip graphs of bounded degree triangulations of a set S of n points in general (non-convex) position in the plane. Like in Section 3.1, let $\mathcal{T}_k(S)$ be the set of triangulations of S such that all vertices have degree at most k. There are point sets for which every triangulation must have a vertex of degree $n - 1$. In other words, $\mathcal{T}_k(S) = \emptyset$ for any $k < n - 1$. Nevertheless we can ask the following question: If there are two triangulations $T_1 \in \mathcal{T}_k(S)$ and $T_2 \in \mathcal{T}_k(S)$, is it possible to flip from T_1 to T_2 while at no time exceeding a vertex degree of k?

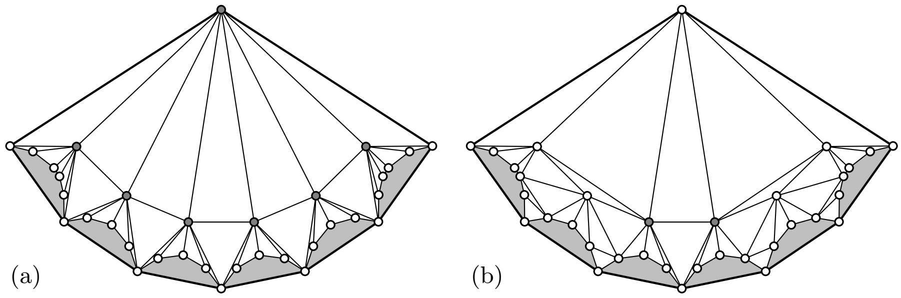

Figure 3.12: Two triangulations which cannot be flipped into each other.

We can answer this question in the negative for $k = O(\sqrt{n})$. For $k < 7$ we know from Section 3.1 that the flip graph of $\mathcal{T}_k(S)$ is not connected. Further, there exists a point set which has two triangulations $T_1 \in \mathcal{T}_k(S)$ and $T_2 \in \mathcal{T}_k(S)$ which cannot be flipped into each other without exceeding a vertex degree of $k \geq 7$. Consider the example for $k = 8$ depicted in Figure 3.12. The shaded parts indicate zigzag triangulations and the dark vertices have degree 8. In the left triangulation, only edges of the zigzags can be flipped without exceeding vertex degree 8. Hence it is impossible to reach the triangulation on the right.

This example can be easily modified for any constant $k \geq 7$. Observe that such an example will have $n = \Theta(k^2)$ points. Thus $k = \Theta(\sqrt{n})$. To get k to be any function of n from constant to square root, we have to be able to add an arbitrary number of points to the example. This can be done inside one of the zigzag areas (shaded parts in Figure 3.12), that has $k - 2$ points on its convex hull. In such a region, there exists a zigzag triangulation with a triangle, $\triangle$, having vertices of degree $(3, 4, 5)$ in both triangulations, T_1 and T_2. It is easy to see that one, two, or three points can be inserted inside $\triangle$, such that there exists a triangulation in which neither the vertices of $\triangle$ nor the inserted points exceed a vertex degree of 6. Furthermore, if three points are inserted, they form another triangle with vertex degrees $(3, 4, 5)$, in which more vertices can be inserted, and so on. With this construction we can build examples of point sets for any k ranging from $k = \Theta(1)$ to $k = \Theta(\sqrt{n})$.

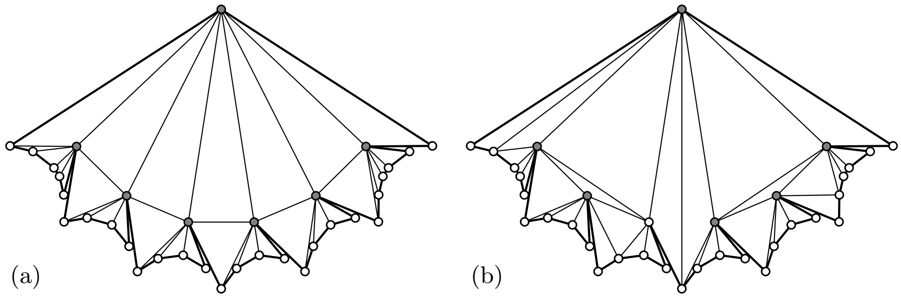

Figure 3.13: Two triangulations of a simple polygon which cannot be flipped into each other.

For triangulations of simple polygons (without inner points) the situation is the same. Using similar arguments as for point sets, we can prove that there exists a simple polygon P such that the flip graph of $\mathcal{T}_k(P)$ is not connected. An example for $k = 8$ is shown in Figure 3.13. This example is obtained by "cutting" the example from Figure 3.12 apart, removing the shaded regions, and slightly adapt the triangulation of Figure 3.12(b).

3.4 Pointed pseudo-triangulations

We briefly consider the problem of degree-bounded flipping for pointed pseudo-triangulations. They are of particular interest among the class of pseudo-triangulations for many reasons. For instance, and contrary to triangulations, in pointed pseudo-triangulations each internal edge can be flipped, and they have exactly $2n-3$ edges, independent of the size of the convex hull of the point set. This is the minimum number of edges a pseudo-triangulation of a given point set can have. Thus pointed pseudo-triangulations are sometimes also called minimum pseudo-triangulations, which is for example the case in [27] and [57] cited below.

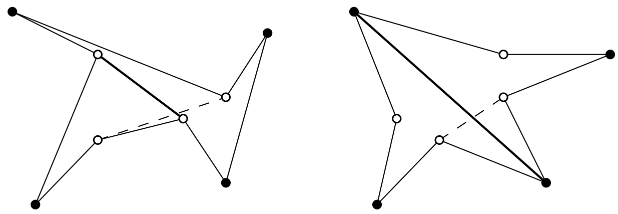

Figure 3.14: Two examples of (exchanging) edge flips in pointed pseudo-triangulations.

Recall that a flip in a triangulation exchanges the diagonal of a convex quadrilateral with its unique counterpart. A flip in a pointed pseudo-triangulation exchanges the diagonal of a pseudo-quadrilateral (a polygon that has exactly four corners) with its unique counterpart. See Figure 3.14 for two examples. An edge flip (also called exchanging flip) in a pointed pseudo-triangulation is constant sized, as it exchanges exactly one edge with another one. It is also reversible in the same sense as for triangulations. But, unlike for triangulations, an edge flip on pointed pseudo-triangulations is not necessarily a local transformation, as the involved pseudo-triangles can have linear size. Still, in spite of the general definition for flips given in the introduction of this chapter, it became common usage to call this operation a flip on pointed pseudo-triangulations. We will be back on flips in triangulations and (pointed) pseudo-triangulations in more detail in Section 4.1.2 of the next Chapter.

The flip graph $\mathcal{F}_{PT}(S)$ of pointed pseudo-triangulations of a point set S is connected and Bereg [27] showed that its diameter is $O(n \log n)$. Further it is known [57] that any point set S in general position has a pointed pseudo-triangulation of maximum vertex degree 5. Hence the question arises if there is a $k \geq 5$ such that the flip graph of pointed pseudo-triangulations with maximum vertex degree k is connected.

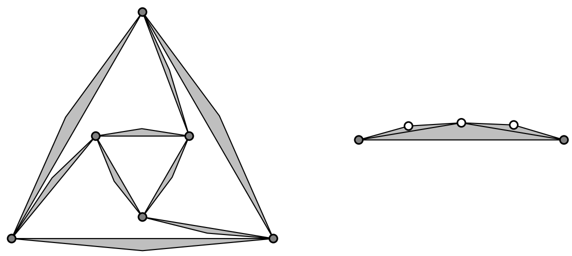

Figure 3.15: Pointed pseudo-triangulation with maximal vertex degree 9 which represents an isolated vertex in the flip graph. The "triangular edges" in the left drawing consist of the structure shown on the right, with the indicated orientation. Dark vertices have degree 9.

For point sets in convex position every (pointed) pseudo-triangulation is in fact a triangulation. Thus our results on triangulations of convex point sets apply also for (pointed) pseudo-triangulations. For point sets in general position we do not know if there is any k such that the flip graph of pointed pseudo-triangulations with maximum vertex degree k is connected, but we do know that k, if it exists, needs to be larger than 9. Consider the pointed pseudo-triangulation PT depicted in Figure 3.15. PT has maximal vertex degree 9, but no edge of PT can be flipped. However, PT is clearly not the only pointed pseudo-triangulation of this point set with maximum vertex degree 9.

3.5 Chapter summary

We considered the flip graph of degree-bounded (pseudo-)triangulations with respect to the maximum vertex degree. For point sets S in convex position we showed that for the set $\mathcal{T}_k(S)$ of triangulations with maximum vertex degree k, the flip graph is disconnected for $k \leq 6$, but connected for any $k > 6$. That is, for any $T_1, T_2 \in \mathcal{T}_k(S)$, $k > 6$, there exists a sequence of flips from T_1 to T_2 such that each intermediate triangulation T_i is also in $\mathcal{T}_k(S)$. Introducing a canonical triangulation for convex point sets, the so called left-most zigzag triangulation, we were able to bound the length of such a sequence with $O(n^2)$.

We strongly believe, that this bound can be improved. On the one hand, because for $k = n-1$ the flip distance is $\Theta(n)$, and on the other hand, because we were already able to show an improved bound on a slightly relaxed setting: When flipping from T_1 to T_2, each intermediate triangulation has to be in $\mathcal{T}_{k+4}(S)$. In this case, the flip distance improves to $O(n \log n)$.

One interesting question is whether this result can be improved further. Another task would be to adapt the proof to the non-relaxed version. The $k+4$ relaxation for intermediate triangulations stems mainly from the intention of keeping the proofs as simple as possible. We believe that a closer investigation will lead to a reduction of the additive term, maybe even to zero. Further, we were able to prove that the flip graph is not connected for triangulations $T \in \mathcal{T}_k(S)$ of point sets S in general position, for any k ranging from $\Theta(1)$ to $\Theta(\sqrt{n})$. But maybe the flip graph can be proved to be connected when allowing a relaxed intermediate degree bound?

A closer look at the results (and their proofs) in Section 3.1 reveals that the degree of each vertex changes monotonically above a degree of 7, when flipping from $T \in \mathcal{T}_k(S)$ to the left-most zigzag triangulation of the point set S in convex position. In this spirit, an interesting variation will be to monotonically decrease the vertex degrees greater than 7 and maintain a degree bound of 7 for all other vertices while flipping from any triangulation to the left-most zigzag triangulation. Or more precise, for a point set S in convex position, let T_j, with $0 \leq j \leq l$, be the sequence of triangulations when flipping from any triangulation of S to the left-most zigzag triangulation T_l of S. We denote by $\delta_j(v_i)$ the degree of vertex v_i in triangulation T_j. We only allow edge flips, such that $\delta_{j+1}(v_i) = max\{7, \delta_j(v_i)\}$ for every vertex v_i. Is the flip graph still connected? Does the $O(n^2)$ bound still hold? Can we improve?

For degree-bounded pointed pseudo-triangulations in general point sets, we showed that for a maximum vertex degree of at most 9 the flip graph is not connected. Hence, for this case, we only know that if the flip graph is connected then the maximum vertex degree has to be at least 10. Apart from the obvious open problem of deciding this instance, also the questions about the above discussed variations apply. For instance, the already mentioned degree bound of 5 for vertices of pointed pseudo-triangulation of general point sets also applies to pointed pseudo-triangulations of (the interior of) simple polygons [14].

Further, for pointed pseudo-triangulations another interesting question arises, when bounding the face degree, that is the maximum number of vertices of pseudo-triangles. Trivially, for convex point sets the flip graph of pointed pseudo-triangulations with maximum face degree 3 is connected. For point sets in general position, there always exists a pointed pseudo-triangulation with maximum face degree 4 [57]. Observe that flipping an edge of a triangle can never result in an increased face degree. As there always exists at least one triangle in a pointed pseudo-triangulation of a point set, there cannot exist isolated vertices in the flip graph of face degree-bounded pointed pseudo-triangulations. This makes it hard to find counterexamples even for small face degree bounds. Thus we ask the following question: Is the flip graph of pointed pseudo-triangulations with maximal face degree 4 connected?

Chapter 4

Pre-triangulations

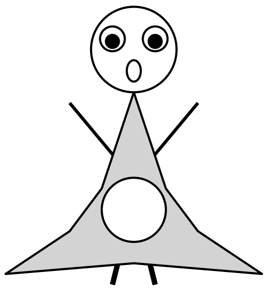

In the previous chapters, we introduced a relaxation of triangulations and a commonly used means of transformation, the flip. In this chapter we will relax a little more and present a new flipping scheme that fits this relaxation.

In Chapter 2 we explored the possible benefits of pseudo-triangulations over triangulations in view of the minimum weight optimality criterion. The potential advantage is due to the possibly reduced number of edges of pseudo-triangulations and was explored to the point of pointed pseudo-triangulations. Thereby we exemplified the relaxation of triangulations to pseudo-triangulations. In the present chapter we will go one step further. We will show that pseudo-triangulations relax to an even more general structure, a structure that turned out to be more natural than pseudo-triangulations in some cases: the so called pre-triangulations.

In Chapter 3, we introduced flips in triangulations and pointed pseudo-triangulations. Due to the presented results, our main focus was on flipping (with certain constraints) in triangulations of convex sets. Nevertheless, the idea of a general transformation scheme was described and we learned the relevance of flip graphs. Further, even if not explicitly stated, we learned that triangulations and also pointed pseudo-triangulations are closed under the edge exchanging flip. That is, flipping an edge in a triangulation always results in another triangulation. And the same is true for edge flips in pointed pseudo-triangulations.

As we will briefly mention in the present chapter, there exists a new type of flip for pseudo-triangulations, the so called edge removing flip. Using this flip and its inverse, the edge inserting flip, it is possible to flip from a triangulation of a point set, S, to a pointed pseudo-triangulation of S and back. Interestingly, it has been shown [10], that allowing edge exchanging and edge removing/inserting flips, the flip distance between two triangulations of a point set is $O(n \log n)$. After introducing pre-triangulations in detail, we will present a new flipping scheme. This new flip is derived from surfaces and is necessary for a "correct" transformation of one pre-triangulation into another.

The results from this chapter have already been presented at [6] and a journal paper was already published in [5].

Note that we relax not only the polygonal complexes but also the underlying domain in the present chapter. Thus it will be necessary to make some already stated definitions stricter. This will not change the general meaning of these definitions, but is necessary for mathematical correctness and compatibility to other definitions. To make it easier for the reader we will also repeat most of the already earlier stated definitions.

The aim of this chapter is to generalize triangulations in a natural way beyond pseudo-triangulations. Stated in slightly other words than before, a pseudo-triangle is a simply connected polygonal region where exactly three vertices have no reflex angle. Dropping simplicity, we arrive at a concept we will call a *pre-triangle*, and following suit, a *pre-triangulation* of a given domain. We show that pre-triangulations arise in three different contexts: In the characterization of complexes that are liftable to three-space in a strong sense, in flip sequences for general polygonal complexes, and as graphs of maximal locally convex functions. Below we give some background on these topics and briefly outline our results.

4.1 Preliminaries

Our work focuses on polygonal complexes, also known as cell complexes. We sometimes also simply use "complexes" for brevity. As these structures are mentioned frequently in this introductory section, we give a brief description. Loosely speaking, a polygonal complex is a partition of a (two-dimensional) region using a set of points and edges, such that each internal edge is adjacent to exactly two different faces. Examples for polygonal complexes are triangulations, quadrangulations, Voronoi diagrams, and also the Schlegel diagram depicted in Figure 4.1. A formal definition of polygonal complexes is given in Section 4.2, termed "polygonal complexes", on pages 37 and the following.

4.1.1 Liftable complexes

The issue of characterizing complexes that can be 'lifted' to space has been frequently investigated in the mathematical literature. A classical theorem of Steinitz [86] implies that, for every complex in the plane whose edge graph is three-connected, there exists a convex 3-polyhedron with isomorphic boundary. The Maxwell-Cremona theorem, see e.g. [36], characterizes polygonal complexes whose edges are exactly the vertical projections of the edges of a polyhedral surface. Complexes with this property are sometimes called *projective* or *liftable* complexes. When the projection surface is required to be convex, the well-studied class of *regular* complexes is obtained. Several criteria for characterizing regular complexes (in general dimensions) exist; see e.g. [18] for a short bibliography. Deciding liftability is also a question of practical importance, for example, in scene analysis.

In the present chapter, we are interested in a concept of liftability which is robust with respect to variations in lifting heights, and, at the same time, outrules complexes that are (not) liftable only because of geometric artifacts. We introduce the notion of *combinatorial projectivity*, which informally means that after a random ε-perturbation of the complex, the space of exact lifts has the maximum possible dimension with probability one. Interestingly, this dimension is related to the number of those vertices in the complex where the angles at all their incident faces are convex. This number we will call the *degree* of the complex (not to be confused with the vertex degree or face degree from the previous chapter). Every polygonal complex has an exact lifting dimension at most its degree, and there are complexes achieving this bound. Combinatorial projectivity is, thus, a property of the graph of edges of the complex and of the order type[1] of its vertices, rather than of the geometry of the vertices and their surface heights.

To see some examples, triangulations are projective complexes that are also combinatorial projective (though they are not necessarily regular). Schlegel diagrams [49], and thus Voronoi diagrams and power diagrams [18], are regular and therefore projective. However, these complexes are not combinatorial projective unless they are triangulations: A movement of vertices of degree three will typically destroy their projectivity, and thus the existence of an exact lifting surface; see Figure 4.1. Also, non-triangular faces will split for most choices of heights for surface vertices. For an example, consider Figure 4.2(right) without the bold edges, which is again a Schlegel

[1]The *order type* [47] of a finite set S of points in the plane assigns to each ordered triple $\{p, q, r\} \subset S$ its orientation, either clockwise or counterclockwise.

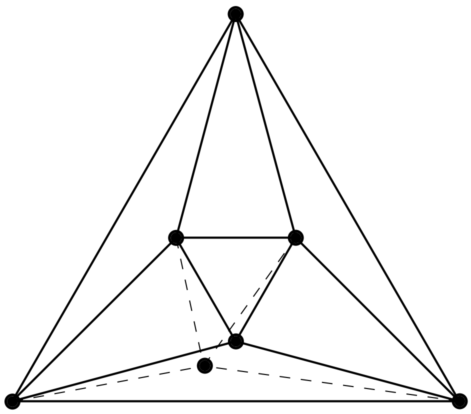

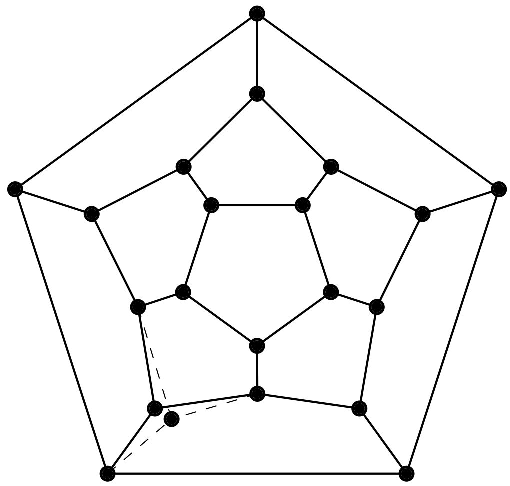

Figure 4.1: Two Schlegel diagrams with perturbed vertex (vertex with dashed lines). (left) A Schlegel diagram which is a triangulation. The perturbation does not destroy the projectivity in this special case. (right) An arbitrary small perturbation destroys the projectivity of this Schlegel diagram.

diagram. A small movement of one vertex will split all quadrilaterals and, as one possibility, introduce the displayed bold edges. Especially, complexes consisting of a single polygonal face are not combinatorial projective, in general. Their space of exact lifts is of dimension only three, even for convex n-gons (where only convex angles occur at vertices, such that the above-mentioned bound on the lifting dimension is n). Clearly, such complexes are projective, but an exact lift exists only for vertex heights with special geometry.

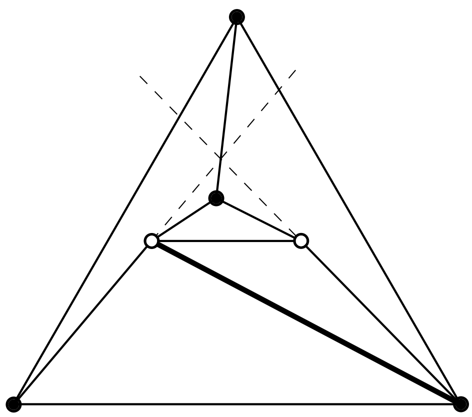

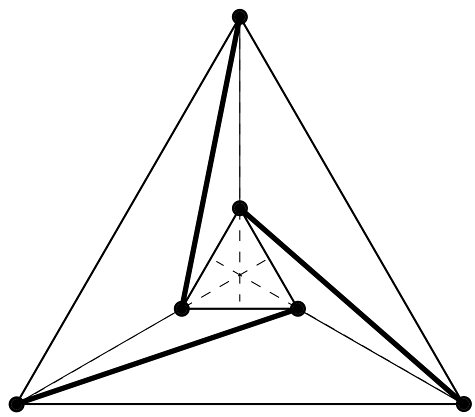

Figure 4.2: Pseudo-triangulations with flat surface edge (left one taken from [10]). Dashed lines are prolongations of edges. Bold edges will not appear in any lifting.

On the other hand, the complexes shown in Figure 4.2 are not projective, because the edges drawn in bold flatten out in all possible projection surfaces: The pairs of triangles sharing these edges always lie in the same plane, which concurs with the two adjacent planes at the point where the dotted lines cross. But ε-perturbing almost surely restores projectivity in these cases. Moreover, these complexes are combinatorial projective, as an exact lifting surface will exist for almost all choices of heights for the vertices where all incident angles are convex (emphasized in black).

In Sections 4.3 through 4.6, we derive a criterion that completely characterizes when a given polygonal complex is combinatorial projective. This characterization hinges on the concept of pre-triangulations, introduced in Section 4.3. More specifically, a complex is combinatorial projective if and only if it is identical to its *M-skeleton*, a certain pre-triangulation defined in Section 4.4. In that section we also show that the surface theorem for pseudo-triangulations in [10] holds in a more general setting. Loosely speaking, this theorem asserts that three (non-trivial) vertex heights per face can be chosen in any given pseudo-triangulation such that each face lifts to planarity. We characterize the class of complexes where the surface theorem applies. This leads to the class of so-called *face-reducible* complexes which contains the pre-triangulations as a proper subclass.

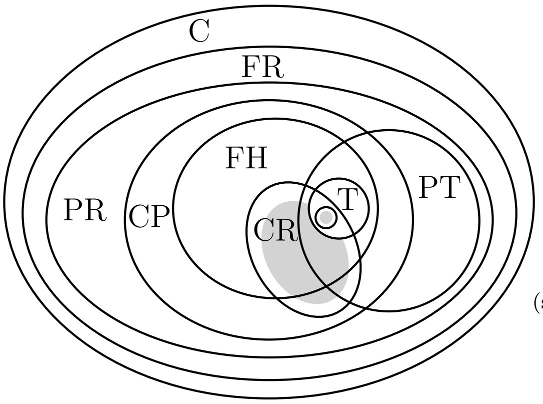

C	... polygonal complexes (page 38)
FR	... face-reducible complexes (page 41)
PR	... pre-triangulations (page 39)
CP	... combinatorial projective complexes (page 44)
FH	... face-honest complexes (page 46)
CR	... combinatorial regular complexes (page 46)
PT	... pseudo-triangulations
T	... triangulations
○	... Delaunay triangulations constrained by R
(shaded)	... area for Delaunay minimum complex of R
	... (for Delaunay minimum complex see page 41)

Figure 4.3: Hierarchy of polygonal complexes in a given region R.

In Section 4.7 we turn to convex projection surfaces and define *combinatorial regularity*, a property stronger than combinatorial projectivity. We introduce the class of *face-honest* complexes, which are basically those where each single face can be lifted to a different plane. Both face-honest complexes and combinatorial regular complexes are certain pre-triangulations. Combinatorial regular complexes in simply connected regions are shown to be face-honest pseudo-triangulations. The various considered complex classes and their containment relations are illustrated in Figure 4.3. The picture is non-redundant in the sense that subclasses are proper and class overlaps are nonempty. Figure 4.4 presents one simple example for each of these 15 cells (not including the Delaunay minimum complexes). At this stage this figure is meant to give the reader a first view at the differences between the complex classes and provide a reference for the reader to consult later, when these complex classes get defined. The same examples will be discussed in more detail in Section 4.13 on page 54 at the end of the chapter.

4.1.2 Flips and convexity

One of the most basic properties of convex sets is their facial structure [31, 49]. In this sense, every convex and piecewise linear function generates a cell complex in its domain of definition. For example, the maximal convex function that does not exceed certain prescribed values on a finite set of points in the plane leads to a complex whose faces are polygonal. This insight dates back to the classical observation that Delaunay triangulations are projected lower convex hulls; see e.g. [21, 43]. A similar relation exists in a more general setting [10]: Maximal *locally convex* functions generate constrained regular pseudo-triangulations if the domain of definition is a nonconvex polygon. We extend the situation to more general domains. In Section 4.8 we prove that graphs of locally convex functions on arbitrary polygonal domains (with possible holes) generate combinatorial regular complexes, and thus pre-triangulations, in the generic case. This result turns out useful in the design of a new and powerful flipping operation for polygonal complexes, whose properties are studied in Sections 4.9 to 4.12.

As already mentioned in Chapter 3, *flips* are a frequently used means to modify polygonal complexes. Various types of flip operations have been considered in the literature, mainly for triangulations and pseudo-triangulations. To give a few relevant citations, we refer to [40, 56, 62] and [10, 71, 75, 88], respectively. All these flips can be defined via geodesic lines in the domain

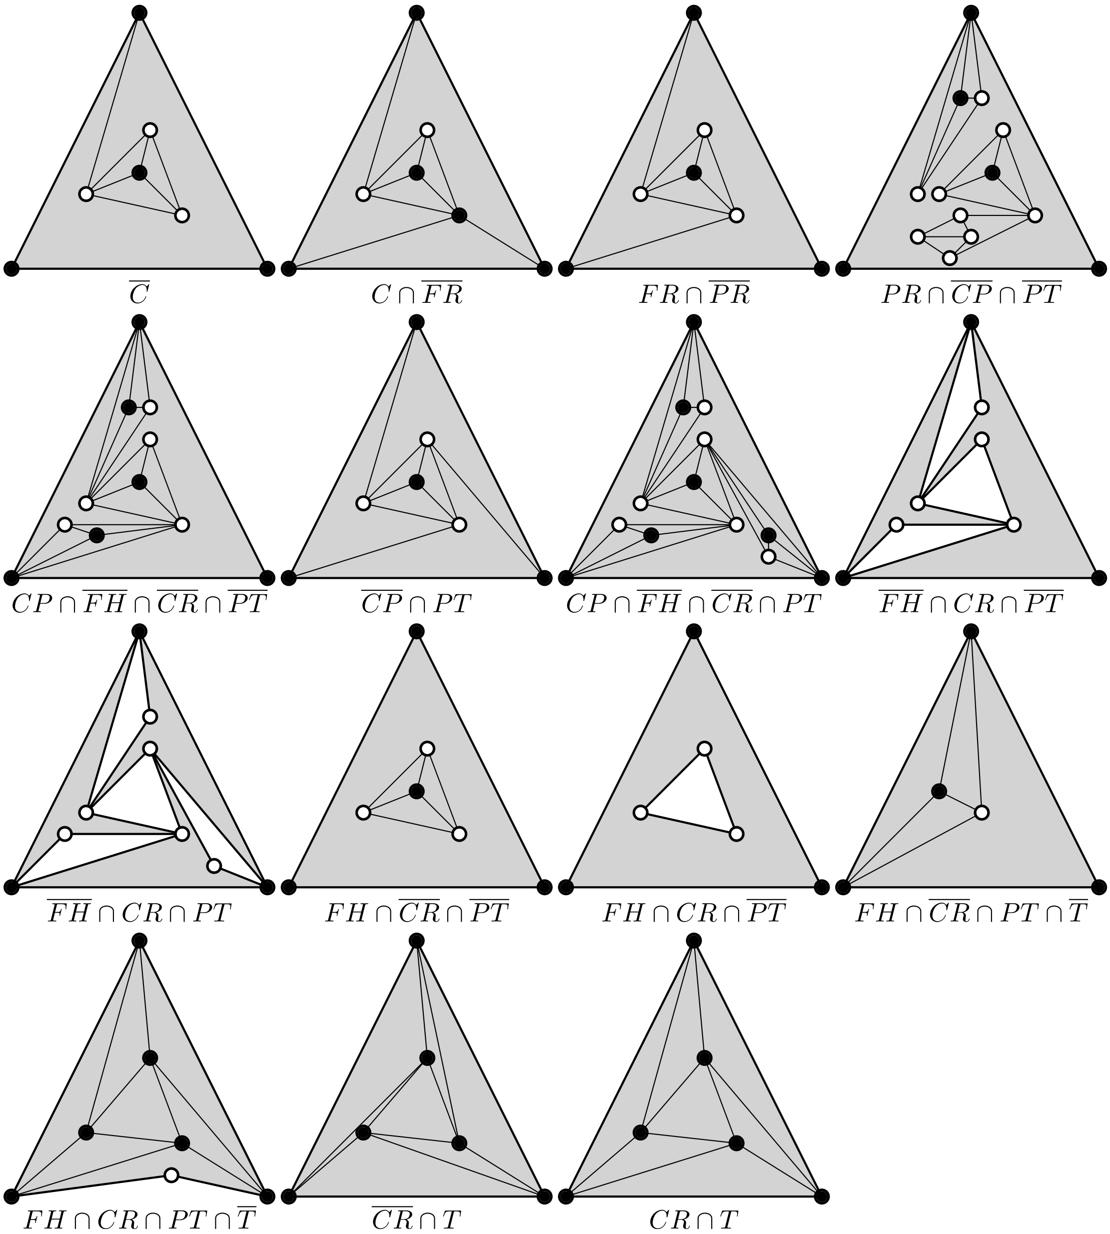

Figure 4.4: Simple examples of polygonal complexes for each cell in the hierarchy of complex classes, shown in Figure 4.3.

where the flip takes place: An edge e to be flipped is adjacent to exactly two pseudo-triangles. In each pseudo-triangle e is part of one of the three reflex chains between two corners. The third corner is the unique corner which is opposite to e. The geodesic line connects the two corners (from both pseudo-triangles) that are opposite to e. Either exactly one straight line segment of this geodesic line is new, or it consists solely of straight line segments which are already part of the boundary of one of the two involved pseudo-triangles. The former case describes an edge exchanging flip, while the latter one describes an edge removing flip. Both types of edge flips (and the inverse of the second one) will be briefly discussed in the following.

Figure 4.5 shows examples for *edge exchanging flips*, the standard Lawson flip for triangulations (left) and the exchanging flip for pseudo-triangulations (middle and right). The edges to be flipped are drawn in bold. The geodesic lines that define the edges created in the flip are shown

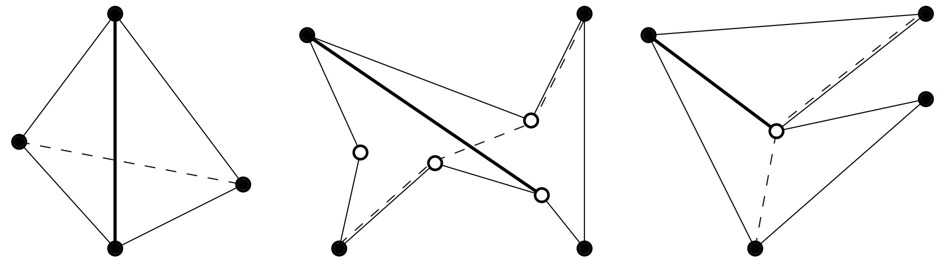

Figure 4.5: Edge exchanging flips.

dashed. The only edge created is the edge of the geodesic line, that does not already exist. Thus, although the geodesic line can be of linear size, the flip is still of constant size, that is, one single edge is exchanged by another one. As already mentioned in Chapter 3, each internal edge of a pointed pseudo-triangulation can be flipped using an edge exchanging flip. This is in general not true for triangulations and pseudo-triangulations. See Figure 4.6, left and middle. Both complexes are pseudo-triangulations, the left complex is a triangulation. In both examples, the bold edge cannot be flipped, as the geodesic line does not introduce any new edge. But if we allow to move between pseudo-triangulations (on the same point set) with different number of pointed vertices these "unflippable" edges become "flippable". This leads to a new type of flip [10].

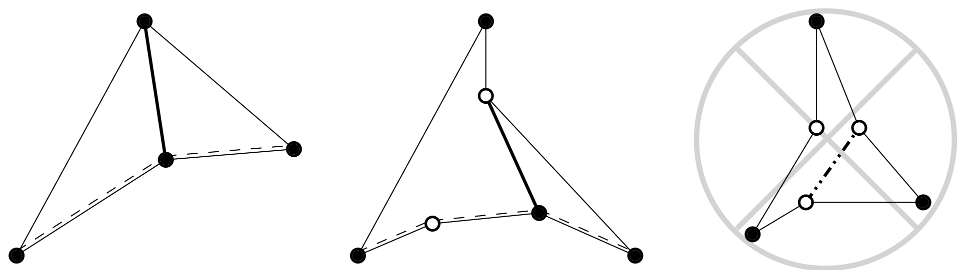

Figure 4.6: Left and middle: Edge removing and edge inserting flips. Right: This edge insertion (dash dot dotted edge) is not a valid flip.

An *edge removing flip*, as illustrated in Figure 4.6 left and middle, removes exactly one edge from a pseudo-triangulation and makes one non-pointed vertex of the polygonal complex pointed. Unlike the edge exchanging flip, the edge removing flip is of course not (self-)reversible. But there exists an inverse operation, the *edge inserting flip*. An edge inserting flip changes the pointedness of exactly one vertex of a pseudo-triangulation from pointed to non-pointed. Loosely speaking, an edge inserting flip can insert an edge, which immediately after its insertion is applicable for an edge removing flip. Figure 4.6(right) exemplifies an invalid edge insertion (dash dot dotted edge). In this case, a pseudo-quadrilateral is created and two vertices of the pseudo-triangulation change their pointedness from pointed to non-pointed.

Using a sequence of edge exchanging, edge removing, and edge inserting flips, it is possible to flip from any pseudo-triangulation of a point set to any other one. As already mentioned at the beginning of this chapter, using these three types of flips, we can bound the flip distance between any two triangulations of a point set by $O(n \log n)$, as shown by Aichholzer et al. [10].

Flipping in liftings of pseudo-triangulations, especially when flipping to the locally convex functions, necessitates another type of flip, not for edges but vertices. Applicable to a pointed vertex with exactly two incident edges, a *vertex removing flip* removes this vertex together with its incident edges from the polygonal complex, see Figure 4.7 middle to right. The pointedness of the remaining vertices remains unchanged. It is easy to see that if a vertex removing flip is applied to a vertex of a pseudo-triangulation, the result is again a pseudo-triangulation. A non-pointed vertex with three incident edges can be removed by applying an edge removing flip to one of the

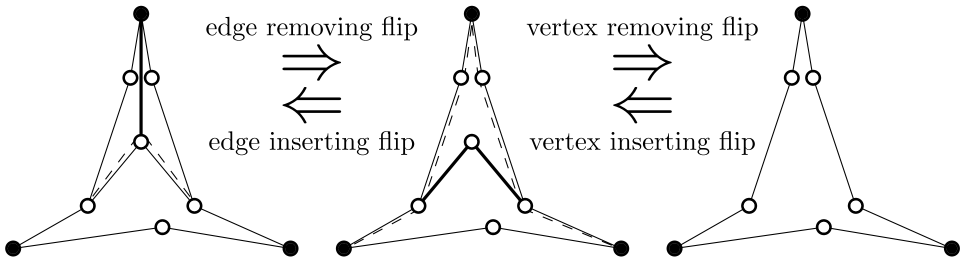

Figure 4.7: Example flip sequence for removing a vertex of degree 3. Left: The edge removing flip results in a pointed vertex with 2 incident edges. Middle: The vertex removing flip deletes the vertex and its incident edges. Right: The resulting polygonal complex.

three edges immediately followed by a vertex removing flip. The inverse operation of a vertex removing flip is the *vertex inserting flip* and consists in principle of the reversed procedure for vertex removing flips. The sequence of flips for removing a non-pointed vertex with three incident edges and the reversed sequence are depicted in Figure 4.7.

The class of pseudo-triangulations is closed under flips of these five types (edge exchanging, edge removing and inserting, vertex removing and inserting). Meaning, if a flip of any of these types is applied to an internal edge of a pseudo-triangulation, the resulting polygonal complex is a pseudo-triangulation again. With this respect, the class of triangulations is closed under edge exchanging flips (the standard Lawson flip).

Extending a different approach, taken in [10] and based on locally convex functions rather than on geodesic lines, we derive a general flip operation in Sections 4.9 and 4.10. This operation works for the entire class of face-reducible complexes and covers all the classical flip types. The smallest class (from the classes listed in Figure 4.3) that is closed under this operation are the pre-triangulations.

In Section 4.11 we introduce the *Delaunay minimum complex*, which is a variant of the well-known Delaunay triangulation [21, 43] for general polygonal domains. This (unique) structure is the complex of smallest combinatorial size that still retains the desired 'Delaunay properties': an analog of local Delaunayhood, and reachability by improving flips. The Delaunay minimum complex is strongly related to the concept of constrained regular pseudo-triangulations introduced in [10], and generalizes the concept of pointed Delaunay pseudo-triangulation in [80]. We prove that any given face-reducible complex in a polygonal region R can be flipped to the Delaunay minimum complex of R by means of improving flips. This connectivity result is a generalization of the optimality theorem in [10]. Section 4.12 shows that, within simply connected regions, every triangulation can be flipped to a predefined constrained regular pseudo-triangulation, in a way such that all intermediate complexes are face-honest pseudo-triangulations (except immediately before a vertex is removed). Flip sequences that retain face-honesty are desirable because they change the complex at hands in a local way.

4.2 Polygonal complexes

This section provides the definitions and notions we will use to work on polygonal complexes. As already mentioned, some definitions from previous chapters will be restated in more detail and in a maybe slightly different style. Nevertheless, the "meaning" of the definitions remains unchanged.

Let R be a bounded subset of the plane. We call R a *polygonal region* if the boundary of R is piecewise linear and coincides with the boundary of the interior of R. Neither connectedness nor simple connectedness of a polygonal region R is required. The boundary components of R are called *edges* and *vertices* of R. A *corner* of R is a vertex of R with no internal angle larger than π. All other vertices of R are termed *noncorners* of R. See Figure 4.8. The depicted region consists of three connected components, two being simply connected. Corners and noncorners are distinguished as black dots and white dots. Note that more than one internal angle may arise at

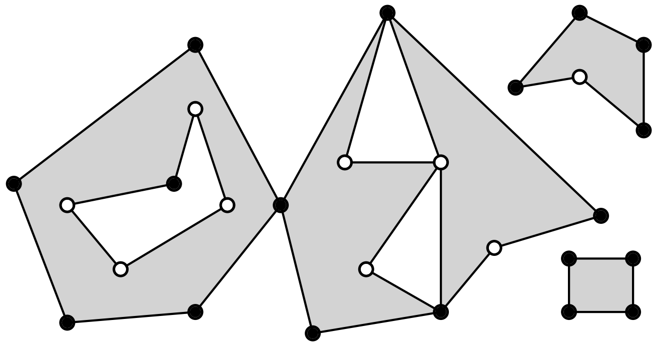

Figure 4.8: General polygonal region.

a single vertex. For corners, all these angles have to be convex.

A (simple) *polygon* is a polygonal region that is homeomorphic to a disk. Observe that the convex hull of a polygonal region R is a convex polygon whose vertices are corners of R. This implies that every polygonal region has at least 3 corners.

A *polygonal partition*, $\mathcal{C}$, of R is a partition of R into (finitely many) polygonal regions. Such a region f is called a *face* of $\mathcal{C}$ if the interior of f is connected. Faces need not be simply connected; they may contain holes and thus are not polygons, in general. The *edges* and *vertices* of $\mathcal{C}$ are the edges and vertices of its faces. An edge (vertex) of $\mathcal{C}$ is called *internal* (to R) if it does not lie on the boundary of R. A polygonal partition $\mathcal{C}$ is termed a *polygonal complex* in R if each internal edge is an edge of two different faces.

We will restrict attention to polygonal complexes in this paper. The vertices of any polygonal complex $\mathcal{C}$ are assumed to be in general position[2] in the plane. As the vertices of the underlying polygonal region R arise as vertices of $\mathcal{C}$, the vertices of R are required to be in general position as well.

Consider an arbitrary subset B of faces of a polygonal complex $\mathcal{C}$. Let v be a vertex of some face in B. Adopting notation from [10], vertex v is called *complete in* B if v is a corner of each face in B that is incident to v. Otherwise, v is called *incomplete in* B. Note that, if v is incomplete in B, then there is a unique face in B where the internal angle at v is larger than π. For vertices of $\mathcal{C}$ which are not incident to some face in B, the completeness status with respect to B is left undefined. Distinguishing complete and incomplete vertices will be crucial in deciding liftability of a complex.

For notational convenience, we will mostly write a polygonal complex $\mathcal{C}$ as the set of its faces. If a vertex v is complete in $\mathcal{C}$, then there is no subset B of faces of $\mathcal{C}$ where v is incomplete. Equivalently, if v is incomplete in $B \subset \mathcal{C}$, then v is also incomplete in $\mathcal{C}$. We define the *degree* of B as the number of vertices of $\mathcal{C}$ that are complete in B. The notion of degree of a face set is central for the developments in the present paper. Notice that the union of the faces in B is a polygonal region, and that each corner of this region has to be complete in B. This implies that the degree of B, and in particular the degree of $\mathcal{C}$, is at least 3.

Figure 4.9 shows a polygonal complex of degree 8. Complete and incomplete vertices are drawn as black dots and white dots, respectively. We will keep this convention (which is also compatible with Figures 4.1, 4.2, and 4.8) throughout this paper. Observe that the (in)completeness of a vertex is a property based on the underlying *complex*, whereas the (non)corner property of a vertex is based on the respective *polygonal region*. Note finally that adjacent faces of a complex may touch at many edges, but each internal edge has to belong to exactly two faces.

[2]A set S of points in the plane is in *general position* if no 3 points of S are collinear.

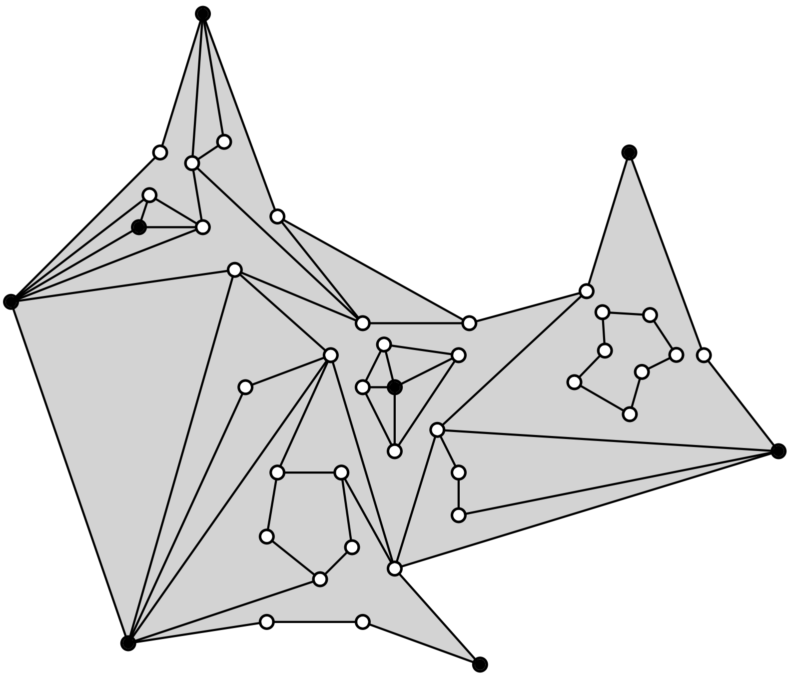

Figure 4.9: A polygonal complex.

4.3 Pre-triangulations

This section introduces the concept of pre-triangulation. We start by recalling the (related) definition of a pseudo-triangulation [10, 71, 75, 87, 88]. We then define minimum complexes in polygonal regions, and finally give a characterization of face sets of degree 3. Such face sets will turn out important in most of our subsequent investigations.

A *pseudo-triangle* is a polygon with exactly 3 corners. A *pseudo-triangulation* is a polygonal complex all whose faces are pseudo-triangles.

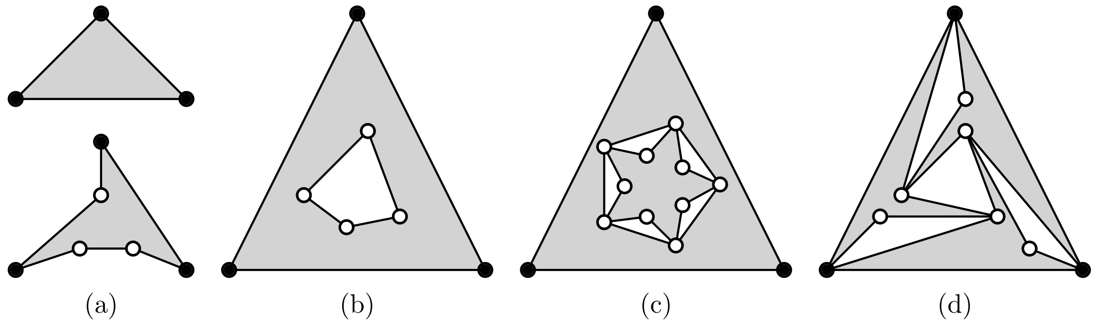

Figure 4.10: Five valid pre-triangles.

We define a *pre-triangle* as an arbitrary polygonal region with exactly 3 corners. Clearly, any pseudo-triangle is a pre-triangle, but the latter may contain convex holes, because no vertex of such a hole is an additional corner. In fact, the interior of a pre-triangle may be disconnected, as it may consist of many edge-disjoint faces. See Figure 4.10. Part (a) illustrates two pseudo-triangles, whereas parts (b), (c), and (d) show pre-triangles which are not pseudo-triangles. The pre-triangle in (c) consists of two faces, and the pre-triangle in (d) consists of three faces.

A polygonal complex $\mathcal{C}$ is called a *pre-triangulation* if $\mathcal{C}$ can be partitioned into subsets

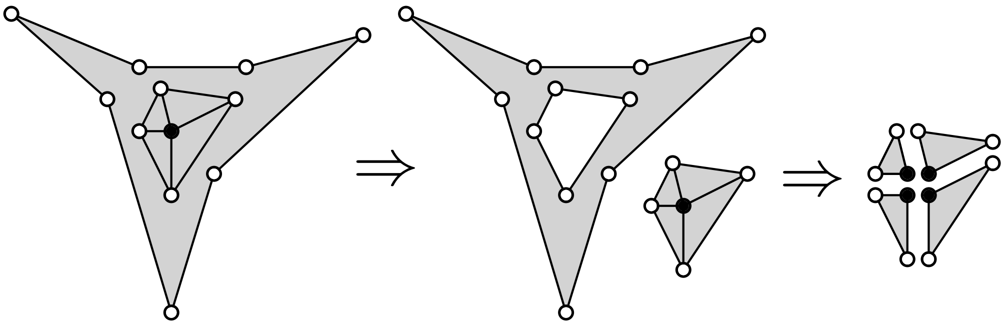

Figure 4.11: The depicted pre-triangulation (left) consists of five pre-triangles: one pre-triangle with a convex hole (middle) and four triangles (right) forming a triangulation inside that hole.

$B_1, \ldots, B_t$ of faces such that (1) the union of the faces in each B_i is a pre-triangle, and (2) the faces of each B_i are pairwise edge-disjoint. Clearly, every pseudo-triangulation is also a pre-triangulation. However, a pre-triangulation may contain faces with holes (Figure 4.11(middle)) and even faces with a large number of corners (Figure 4.10c). Figure 4.11 shows a small example of a pre-triangulation (left) consisting of five pre-triangles. Four of which are triangles forming a triangulation (right). This triangulation results in a convex hole of the fifth face (middle, showing also the triangulation). This exemplifies that holes in pre-triangles need not stem from holes in the underlying polygonal region. Note that the depicted pre-triangulation (left) is a subset of a bigger pre-triangulation shown in Figure 4.13 on page 42.

Let $\mathcal{C}$ be some polygonal complex. Recall that each corner of the underlying polygonal region R is complete in $\mathcal{C}$. This leads us to term $\mathcal{C}$ a *minimum* polygonal complex if the corners of R are the only vertices that are complete in $\mathcal{C}$. In other words, the degree of $\mathcal{C}$ equals the number of corners of R, the minimum that can be achieved. Observe that, in a minimum complex, every internal vertex (if any) is reflex in one of its incident faces. To give examples, the complex in Figure 4.12 and also the complex in Figure 4.19 on page 51 are minimum, whereas the complexes in Figures 4.9 on page 39 and 4.13 on page 42 are not.

For pseudo-triangulations, the definition above is consistent with the original definition that uses vertex pointedness: In a pointed pseudo-triangulation [87, 88] (recall that this is often also called a minimum pseudo-triangulation) each (internal) vertex is pointed, that is, its incident edges span a convex angle. Note that, as we consider general polygonal domains R, a vertex of R can already be non-pointed. Obviously such a vertex will also be non-pointed in the "pointed" pseudo-triangulation (as in any other polygonal complex of R). Thus using the term "minimum pseudo-triangulation", also with regard that it is the pseudo-triangulation with the fewest edges, makes more sense in this context. Observe that a polygonal complex $\mathcal{C}$ in R is minimum if and only if $\mathcal{C}$ can be 'filled up' to a minimum pseudo-triangulation in R by adding edges between vertices of $\mathcal{C}$. In particular, the number of edges of a minimum complex in R is not determined by its vertex set. The natural counterpart to minimum complexes are complexes where all vertices are complete. Triangulations and Schlegel diagrams (Figure 4.1) are notable examples.

The following lemma characterizes face sets of degree 3 in a polygonal complex $\mathcal{C}$. Here and in later sections we denote with $U(B)$ the union of the faces in a subset $B \subset \mathcal{C}$.

Lemma 15. *Let B be any subset of faces of $\mathcal{C}$. Then B is of degree 3 if and only if $U(B)$ is a pre-triangle and B forms a minimum polygonal complex in $U(B)$.*

Proof. Assume that B is of degree 3. As $U(B)$ has at least 3 corners, and each such corner is complete in B, we conclude that $U(B)$ has exactly 3 corners. That is, $U(B)$ is a pre-triangle, and the polygonal complex formed by B in $U(B)$ is minimum.

Conversely, assume that $U(B)$ is a pre-triangle, and that the polygonal complex formed by B in $U(B)$ is minimum. Then the degree of this complex equals the number of corners of $U(B)$, which is 3 because $U(B)$ is a pre-triangle. That is, B is of degree 3. □

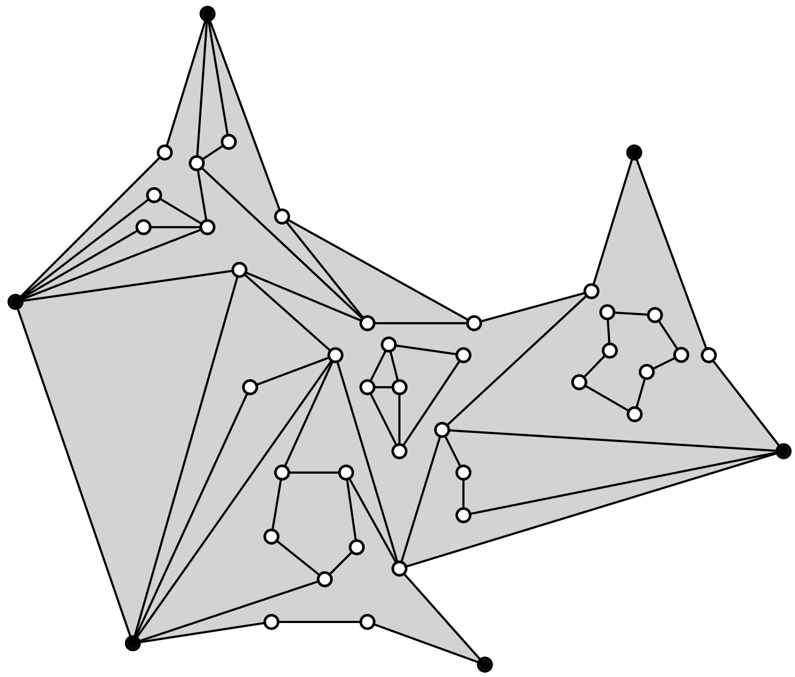

Figure 4.12: Minimum polygonal complex.

4.4 The M-skeleton

Utilizing pre-triangulations, we now define a substructure for polygonal complexes, the so-called M-skeleton, which is the key to combinatorial projectivity. Intuitively speaking, the M-skeleton delimitates maximal face sets of the complex that lift to planarity in a robust sense. We introduce the class of face-reducible complexes and demonstrate that the surface theorem for pseudo-triangulations in [10] can be extended to this more general class. Face-reducible complexes are also the largest class where this is possible, as Section 4.6 shows.

A polygonal complex is termed *face-reducible* if and only if each of its faces is contained in some subset of faces of degree 3. For example, every pre-triangulation (and, in particular, every (pseudo-)triangulation) is a face-reducible complex. Note that such complexes may contain faces of any shape. Figure 4.9 om page 39 gives an illustration. The reader is encouraged to check that this complex is face-reducible (using Figure 4.13 on page 42 as an aid).

Let $\mathcal{C}$ be a face-reducible complex for the rest of this section. For a face $f \in \mathcal{C}$, we will denote with M_f a maximal subset of faces of $\mathcal{C}$ that contains f and that is of degree 3. Recall from Lemma 15 that $U(M_f)$ is a pre-triangle. We state and prove two basic properties of such maximal face sets.

Lemma 16. *Let $f, g \in \mathcal{C}$. If $M_f \cap M_g \neq \emptyset$ then $M_f = M_g$.*

Proof. Let $M_f \cap M_g \neq \emptyset$. We claim that $M_f \cup M_g$ is of degree 3. This implies $M_f = M_g$ by the maximality of these sets.

Put $Q = U(M_f \cup M_g)$ and $I = U(M_f \cap M_g)$. Using counting arguments for the vertices of U and I, we prove that Q has exactly three corners, and that $M_f \cup M_g$ forms a minimum complex in Q. This, by Lemma 15, then implies that $M_f \cup M_g$ is of degree 3.

The sets M_f and M_g form minimum complexes in $U(M_f)$ and $U(M_g)$, respectively (Lemma 15). Hence, a vertex is complete in M_f if and only if it is a corner of $U(M_f)$ (same for M_g). Moreover, a vertex v being incomplete in M_f or M_g is incident to a unique face $\mu \in M_f \cup M_g$ at a reflex angle, and if v is incomplete in both M_f and M_g then $\mu \in M_f \cap M_g$ holds. We thus observe two properties:

(a) If c is a corner of Q *or* of I then c is a corner of at least one of $U(M_f)$ and $U(M_g)$.

(b) If c is a corner of Q *and* of I then c is a corner of both $U(M_f)$ and $U(M_g)$.

Now, $U(M_f)$ and $U(M_g)$ together have 6 corners, and I has at least 3 corners. So, by properties (a) and (b), Q has at most (and thus exactly) 3 corners.

Concerning the complex in Q, we recall that M_f and M_g already form minimum complexes. Moreover, any vertex being incomplete in M_f or M_g has to be incomplete in $M_f \cup M_g$ as well. So the minimum property of the complex in Q is violated only if a vertex, w, that is complete in both sets M_f and M_g becomes a noncorner of Q. Observe that w has to be a corner of both $U(M_f)$ and $U(M_g)$ then. So, from the pool of 6 corners we get from $U(M_f)$ and $U(M_g)$, already 2 are used for the noncorner w of Q. Another 3 are used for corners of Q, by (a). To cover the corners of I according to (a), 1 corner might be w, and 2 additional ones can be either yet unused corners, or corners of Q. In either case, by (b), 2 more corners from the pool are needed, resulting in a total of 7 corners – a contradiction. We conclude that the vertex w does not exist. □

Corollary 3. *The set M_f is unique for each face $f \in \mathcal{C}$. Moreover, the collection of these sets defines a partition of $\mathcal{C}$.*

Proof. As we have $f \in M_f$ for each set M_f, the uniqueness of M_f follows from Lemma 16. To prove the partition property, let us write $f \sim g$ if $M_f = M_g$ holds for $f, g \in \mathcal{C}$. Then $\sim$ is an equivalence relation on $\mathcal{C}$. We show that $[f]_\sim = M_f$ are its equivalence classes. Clearly, $g \in [f]_\sim$ implies $M_g = M_f$ and thus $g \in M_f$. Conversely, $g \in M_f$ implies $g \in M_g \cap M_f$ and thus $M_g = M_f$ by Lemma 16, that is, $g \in [f]_\sim$ holds. □

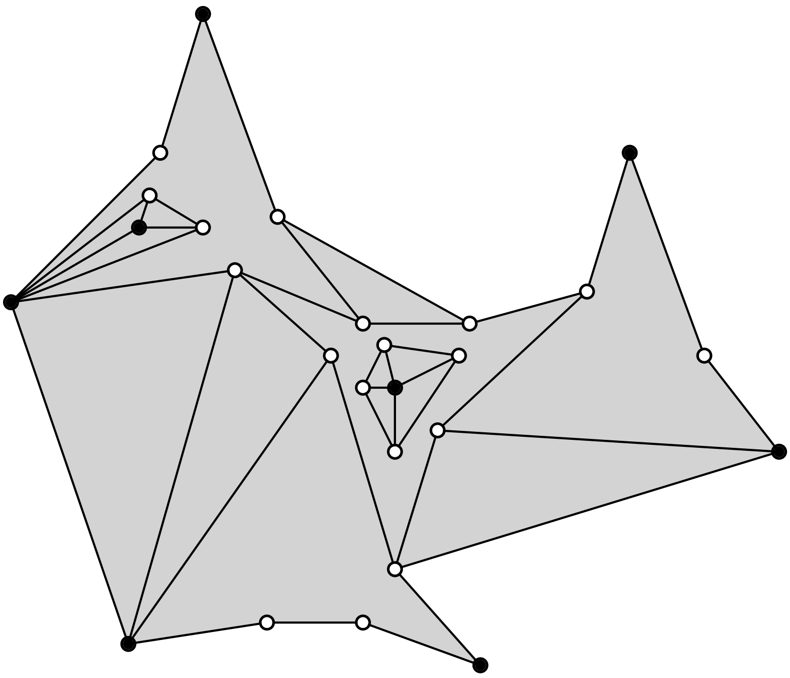

Figure 4.13: M-skeleton of the face-reducible complex in Figure 4.9.

From Corollary 3 and Lemma 15 we know that the regions $U(M_f)$, for all $f \in \mathcal{C}$, partition the underlying region $U(\mathcal{C})$ of $\mathcal{C}$ into pre-triangles in a unique way. We will term the resulting pre-triangulation the *M-skeleton* of $\mathcal{C}$. Figure 4.13 illustrates an example. Note that not all faces of this pre-triangulation are simply connected, although there are no holes in the underlying region.

The following lemma is similar in spirit to Lemma 16. It will be needed in the design of a general flipping operation in Section 4.9.

Lemma 17. *Let $f, g \in \mathcal{C}$ be two faces that share some edge. Then either $M_f = M_g$ or the degree of $M_f \cup M_g$ is 4.*

Proof. Assume $M_f \neq M_g$. Then M_f and M_g are disjoint, by Lemma 16. $U(M_f)$ and $U(M_g)$ have an edge e in common, because f and g do so, and $e = U(M_f) \cap U(M_g)$ holds because we have $M_f = M_g$, otherwise. Like in the counting argument in the proof of Lemma 16 (put $I = e$), each endpoint c of e leads to a loss of at least one of the 6 possible corners of $Q = U(M_f \cup M_g)$. Thus Q has at most 4 corners. Moreover, if an endpoint c of e is complete in $M_f \cup M_g$ but is a noncorner of Q, then c causes a loss of two possible corners of Q. We conclude that $M_f \cup M_g$ is of degree 4. □

Face-reducible complexes enjoy a strong lifting property which stems from the existence of their M-skeletons, and that we are going to describe next. Define a *polyhedral surface* as (the graph of) a continuous and piecewise-linear function φ whose domain is a polygonal region R. $\varphi|_L$ is called a *facet* of φ if L is a maximal interior-connected subset of R where φ is linear.

Theorem 4. *Let $\mathcal{C}$ be a face-reducible polygonal complex. Further, let η be any vector assigning a height η_i to each complete vertex v_i of $\mathcal{C}$. There exists a unique polyhedral surface φ for $\mathcal{C}$ and η, with $\varphi(v_i) = \eta_i$ for all i, and such that $\varphi|_f$ is a subset of a facet of φ, for all faces $f \in \mathcal{C}$.*

Proof. As $\mathcal{C}$ is face-reducible, the M-skeleton of $\mathcal{C}$ exists (and is unique). Let PR denote this pre-triangulation. The conditions required for the surface theorem in [10] to hold, can be formulated as follows: For each incomplete vertex v of $\mathcal{C}$, there is a unique pre-triangle $\nabla = U(M_f)$ of PR such that (1) v is incomplete in M_f and (2) v lies in the convex hull of the 3 corners of ∇.

Consider condition (1). If v is not a vertex of PR then there is a unique ∇ such that v is internal to ∇. Otherwise, there is a unique ∇ where v is a noncorner, because an internal angle larger than π occurs in ∇ at v. In both cases, v is incomplete in the face set M_f with $\nabla = U(M_f)$, by Lemma 15. Condition (2) holds because the vertices of the convex hull of ∇ are corners of ∇. □

4.5 Planar face sets

Throughout this section, let $\mathcal{C}$ be a face-reducible complex. Theorem 4 makes explicit that $\mathcal{C}$ can be lifted, in various ways, to a polyhedral surface in three-space without introducing new edges. However, not all edges of $\mathcal{C}$ might have their counterparts in this surface, for several reasons. For instance, the choice of the height vector η may force more than three complete vertices to be coplanar in the lifting. Also, geometric degeneracies of $\mathcal{C}$ may be the reason; consult Figure 4.2. Using our concept of M-skeleton, we are able to characterize those edges of $\mathcal{C}$ which, even under 'generic' conditions, cannot be made to show up in the surface.

Let S be the set of all vertices of $\mathcal{C}$. Let S_ε be some replacement within distance ε of each vertex in S, for arbitrarily small $\varepsilon > 0$. (We assumed S to be in general position, so the order type of S_ε equals the order type of S.) An *ε-perturbation*, $\mathcal{C}_\varepsilon$, of $\mathcal{C}$ is the polygonal complex with vertex set S_ε and with the same combinatorial structure as $\mathcal{C}$. As order types are preserved, a vertex is complete in $\mathcal{C}_\varepsilon$ if and only if its original is complete in $\mathcal{C}$.

A subset B of faces of $\mathcal{C}$ is called *combinatorial planar* if, for all ε-perturbations $\mathcal{C}_\varepsilon$ of $\mathcal{C}$, and for all height vectors η for the complete vertices of $\mathcal{C}_\varepsilon$, the subset $\varphi|_B$ of the surface φ for $\mathcal{C}_\varepsilon$ and η lies in a single plane. Recall that φ uniquely exists by Theorem 4.

Lemma 18. *The maximal combinatorial planar face sets in a face-reducible complex $\mathcal{C}$ are the pre-triangles $U(M_f)$, $f \in \mathcal{C}$, of its M-skeleton.*

Proof. Consider such a pre-triangle $U(M_f)$. By definition of an M-skeleton, M_f is of degree 3. So exactly 3 vertices are complete in M_f. For every height vector η for $\mathcal{C}$, the plane through the corresponding 3 points in space determines a possible surface for M_f, because all other vertices of the complex formed by M_f are incomplete in M_f, and thus are incomplete in $\mathcal{C}$. By the uniqueness of this surface, all surface vertices for M_f have to lie in this plane. This reasoning remains true for any ε-perturbation $\mathcal{C}_\varepsilon$ of $\mathcal{C}$. Therefore M_f is combinatorial planar.

It remains to show that M_f is maximal under the condition of being combinatorial planar. Consider the linear system that, given $\mathcal{C}$ and η, describes the unique surface according to Theorem 4; see [10]. The variables in this system are the heights for all the vertices of $\mathcal{C}$, including the incomplete ones. Let now B be any proper superset of M_f. By the structure of the system, the heights of the vertices in B only depend on heights of vertices for supersets B' of B. Moreover, the heights for any such B' can be described by using only the heights for the vertices that are complete in B' (Theorem 4 applied to B'). On the other hand, the face set $M_f \subset B$, by definition, is maximal under the condition of having 3 complete vertices. That is, *every* superset $B' \supseteq B$ has degree at least 4. Therefore, when using the equations of the system, the vertex heights for B cannot be described by only 3 of the system variables (unless the system is overdetermined because of geometric artifacts; these can be removed by ε-perturbing both $\mathcal{C}$ and η). Moreover, no other system with the same solution space can describe the heights for B with fewer variables, because our system takes into account *all* the vertex heights. But a plane in three-space is determined by 3 parameters, and we conclude that there exist ε-perturbations $\mathcal{C}_\varepsilon$ of $\mathcal{C}$ and η_ε of η such that the subset $\varphi|_B$ of the surface φ for $\mathcal{C}_\varepsilon$ and η_ε is not coplanar. That is, B is not combinatorial planar. □

4.6 Combinatorial projectivity

We now define a notion of projectivity (i.e., liftability) for polygonal complexes which is robust with respect to variations in lifting height and, at the same time, outrules complexes that are projective (or non-projective) only because of geometric artifacts.

Consider an arbitrary polygonal complex $\mathcal{C}$. Define an *exact lift* of $\mathcal{C}$ as a polygonal surface whose set of edges projects exactly to the set of edges of $\mathcal{C}$. In an exact lift of $\mathcal{C}$, no edge of $\mathcal{C}$ is allowed to flatten out, and no face of $\mathcal{C}$ is allowed to fold at new edges. Not every polygonal complex admits an exact lift; perturbed Schlegel diagrams (Figure 4.1(right)) are an example.

We are interested in complexes whose space of exact lifts has maximal dimension. It will turn out that this dimension is bounded from above by the degree, k, of $\mathcal{C}$. Let us call the complex $\mathcal{C}$ *combinatorial projective* if, for a random[3] ε-perturbation $\mathcal{C}_r$ of $\mathcal{C}$, the space of exact lifts of $\mathcal{C}_r$ has dimension k with probability 1. Intuitively speaking, for almost all perturbations of a combinatorial projective complex, almost all height vectors for its k complete vertices will lead to an exact lift.

Our aim is to characterize the class of combinatorial projective complexes. The next assertion formulates a main observation in this context. Its proof is immediate from the proof of Lemma 18.

Corollary 4. *Let $\mathcal{C}$ be a face-reducible complex, and let $\eta_r \in [0,1]^k$ be a random height vector for the k complete vertices of $\mathcal{C}_r$. Then, with probability 1, the edges in the surface for $\mathcal{C}_r$ and η_r bijectively correspond to the edges of the M-skeleton of $\mathcal{C}$.*

Theorem 5. *A polygonal complex $\mathcal{C}$ is combinatorial projective if and only if the M-skeleton of $\mathcal{C}$ exists and is identical to $\mathcal{C}$.*

Proof. For the class of face-reducible complexes, the assertion is true by Corollary 4. To complete the proof, we assume that $\mathcal{C}$ is not face-reducible (i.e., the M-skeleton of $\mathcal{C}$ does not exist) and argue that $\mathcal{C}$ cannot be combinatorial projective in this case.

Faces of $\mathcal{C}$ not being part of some set of degree 3 ($\mathcal{C}$ is supposed to contain at least one such face) have more than 3 corners. So, for each such face f, some internal edge can be added to split f into two faces, without changing any vertex from incomplete to complete. We keep adding such edges until a face-reducible complex $\mathcal{C}'$ is obtained. Let $B(f)$ denote the set of faces of $\mathcal{C}'$ that a face f of $\mathcal{C}$ splits into. Now, $\mathcal{C}$ has to have some face f whose set $B(f)$ is not contained in a single pre-triangle of the M-skeleton of $\mathcal{C}'$, as $\mathcal{C}$ would be face-reducible, otherwise. Therefore, in the surface φ' for $\mathcal{C}'_r$ and η_r, the set $\varphi'|_{B(f)}$ is not part of a single facet of φ' with probability 1; see Corollary 4. By the uniqueness of the surface theorem, there is no other way of constructing

[3] drawn uniformly for each vertex from its surrounding ε-disk

the required surface for $\mathcal{C}_r$ and η_r. This implies that, with probability 1, the space of exact lifts for $\mathcal{C}_r$ is strictly less in dimension than the degree of $\mathcal{C}_r$. That is to say, $\mathcal{C}$ is not combinatorial projective. □

Corollary 5. *Every combinatorial projective complex is a pre-triangulation.*

The proof of Theorem 5 also shows the following: Every polygonal complex $\mathcal{C}$ can be refined to a face-reducible one without changing its degree k. Therefore, the space $\mathcal{L}$ of 'face-embedding' lifts of $\mathcal{C}$, in the sense of Theorem 4, is of dimension at most k, and this bound is achieved if $\mathcal{C}$ is face-reducible. The space $\mathcal{E}$ of exact lifts of $\mathcal{C}$ is the complement of a finite union of linear subspaces in $\mathcal{L}$, each subspace representing the lifts that flatten out a fixed edge of $\mathcal{C}$. Thus $\mathcal{E}$ is either empty or the dimension of $\mathcal{E}$ is k as well.

If a complex $\mathcal{C}$ is not face-reducible then — apart from special cases caused by geometric degeneracies — for almost all choices of heights for $\mathcal{C}$'s complete vertices, certain faces of $\mathcal{C}$ will split, in every surface which respects these heights. We thus have:

Corollary 6. *Theorem 4 cannot be extended beyond the class of face-reducible polygonal complexes.*

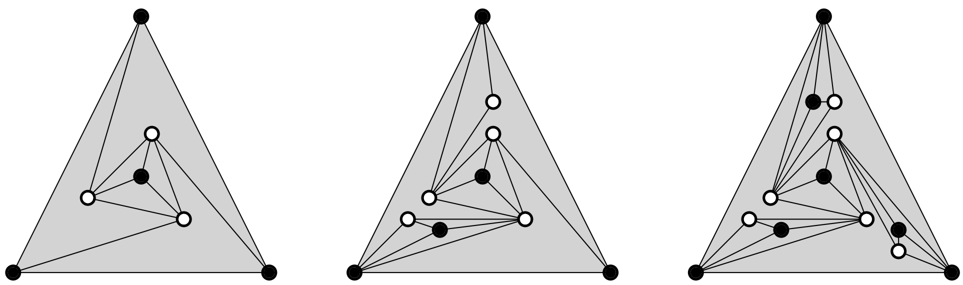

Figure 4.14: Example for combinatorial projectivity.

Figure 4.9 on page 39 shows an example for a face-reducible complex which is not combinatorial projective; the complex is not identical to its M-skeleton (see Figure 4.13 on page 42). A pseudo-triangulation which is not combinatorial projective is shown in Figure 4.14(left): the three pseudo-triangles, which are not triangles, share edges and form a face set which is combinatorial planar. The same is true for four pseudo-triangles (one of them being a triangle) in Figure 4.14(middle). However, combinatorial planarity of a set of two or more faces does not necessarily destroy combinatorial projectivity. That is the case, if the faces of the combinatorial planar face set do not share any edge. For example, the pseudo-triangulation in Figure 4.14(right) constitutes both a combinatorial projective complex and a combinatorial planar face set of size three. Other examples are the pre-triangulations in Figure 4.10 c and d on page 39, which are also combinatorial projective complexes containing combinatorial planar face sets. Note that, although each combinatorial projective complex is a pre-triangulation, not all of its *faces* need to be pre-triangles; see the pentagonal star in Figure 4.10c on page 39.

4.7 Combinatorial regularity

In this section we focus on convex projection surfaces. We introduce the classes of combinatorial regular complexes and face-honest complexes and elaborate on their interrelation. Face-honest complexes are of interest for two more reasons: Loosely speaking, each such complex can be lifted to a surface so that its faces live in pairwise different planes. Moreover, these complexes allow for easy modification with flipping operations; see Section 4.12.

Let φ be some polyhedral surface. An edge e of φ is called *convex* if there exists a line segment ℓ that nowhere lies above φ and that intersects e at exactly one point interior to both e and ℓ. Edges of φ which are not convex are called *reflex*. Note that the image $\varphi|_{e'}$ of each boundary edge e' of the domain of φ is a convex edge.

A polygonal complex $\mathcal{C}$ is called *combinatorial regular* if $\mathcal{C}$ is combinatorial projective and there exists some ε-perturbation of $\mathcal{C}$ that admits an exact lift where all edges are convex. By Corollary 5, combinatorial regular complexes are pre-triangulations. Known examples are (constrained) Delaunay triangulations [66], and more generally, the constrained regular pseudo-triangulations [10]. Schlegel diagrams [49] and thus Voronoi diagrams are not combinatorial regular (though they are well known to be regular in the classical sense) because these complexes are not combinatorial projective, apart from the special case of triangulations.

Lemma 19. *Let $\mathcal{C}$ be a combinatorial regular complex. Then each internal vertex of $\mathcal{C}$ is complete.*

Proof. Let $\mathcal{C}_\varepsilon$ and η define an exact lift that witnesses the combinatorial regularity of $\mathcal{C}$. Consider an incomplete vertex v of $\mathcal{C}$, and let f be the unique face of $\mathcal{C}$ where v is a noncorner. To get a contradiction, suppose v is an internal vertex of $\mathcal{C}$. As there is an internal angle greater than π in f at v, there exists some line segment $\ell \subset U(\mathcal{C})$ such that ℓ crosses all the edges of $\mathcal{C}$ incident to v. But $\mathcal{C}$ is combinatorial regular, so the corresponding edges in the surface φ for $\mathcal{C}_\varepsilon$ and η are all convex. Therefore $\varphi(f)$ cannot be part of a single facet of φ – a contradiction to the definition of φ. □

A (face-reducible) polygonal complex is termed *face-honest* if $M_f = \{f\}$ holds for each of its faces f. Being identical to its M-skeleton, every face-honest complex $\mathcal{C}$ is a pre-triangulation, that is combinatorial projective by Theorem 5. Moreover, all pre-triangles of $\mathcal{C}$ are single faces. Figure 4.13 on page 42 gives an illustration. On the other hand, even when a pseudo-triangulation is combinatorial regular, it is not necessarily face-honest. Figures 4.10 c and d on page 39 reveal this fact. Triangulations are always face-honest: All their vertices are complete, so their individual triangles constitute the only possible face sets of degree 3. See Figure 4.3 on page 34 for the interaction of face-honest complexes and combinatorial regular complexes with other classes.

Below we derive some results for the case where the underlying region $U(\mathcal{C})$ of the complex $\mathcal{C}$ in question is a polygon.

Lemma 20. *Let $\mathcal{C}$ be a face-reducible complex, and suppose each internal vertex of $\mathcal{C}$ is complete. If $U(\mathcal{C})$ is a polygon then $\mathcal{C}$ is a face-honest pseudo-triangulation.*

Proof. Let f be a face of $\mathcal{C}$. As $\mathcal{C}$ is face-reducible, the set M_f exists. Each vertex v that is incomplete in M_f is also incomplete in $\mathcal{C}$, and thus v is a vertex of $U(\mathcal{C})$, by assumption. Therefore, if $U(\mathcal{C})$ is required to be a polygon, then the pre-triangle $U(M_f)$ is just a pseudo-triangle, and M_f contains a single face, $U(M_f) = f$. □

Lemmas 19 and 20 combine to the following statement.

Theorem 6. *Let $\mathcal{C}$ be a combinatorial regular complex whose underlying region is a polygon. Then $\mathcal{C}$ is a face-honest pseudo-triangulation.*

Corollary 7. *Let PT be a pseudo-triangulation without internal vertices in a polygon R. Then PT is face-honest.*

Proof. As PT contains no internal vertices, each internal edge of PT is a diagonal of R. This implies that PT is combinatorial regular. The assertion now follows from Theorem 6. □

4.8 Locally convex surfaces

We show next that combinatorial regular complexes – and thus certain pre-triangulations – arise from graphs of locally convex functions on polygonal domains. This generalizes results in [10] where locally convex functions are introduced and utilized in the context of pseudo-triangulations. The relationship between locally convex functions and pseudo-simplicial complexes in higher dimensions is discussed in detail in [25].

Let a polygonal domain R be given. A surface φ on R is called *locally convex* if each edge of φ is convex.[4] Let $S \subset R$ be a finite set of points that includes all the vertices of R. Further, let η be a vector that assigns an upper height bound $\eta(v)$ to each point $v \in S$. We define F_η as the maximal (i.e., highest) locally convex surface on R that satisfies $F_\eta(v) \leq \eta(v)$ for all $v \in S$.

The surface F_η exists and is unique: Any surface φ that consists of a single facet, $\varphi|_R$, is locally convex and can be lowered to satisfy η, and F_η is the pointwise maximum of all possible surfaces on R that are locally convex and satisfy η. The facets of F_η project to the faces of a polygonal complex in R which we denote by $\mathcal{C}(F_\eta)$.

The next lemma follows from results in [25]. We give a more direct proof here.

Lemma 21. *All vertices of $\mathcal{C}(F_\eta)$ are included in S. For each complete vertex v of $\mathcal{C}(F_\eta)$, $F_\eta(v) = \eta(v)$ holds.*

Proof. Let w be a vertex of $\mathcal{C}(F_\eta)$. Assume $w \notin S$. Then w has to be complete in $\mathcal{C}(F_\eta)$: If w lies on the boundary of R then clearly no face of $\mathcal{C}(F_\eta)$ has an internal angle larger than π at w. If w is an internal vertex of $\mathcal{C}(F_\eta)$ then the same is true, because the image $F_\eta(e)$ of each edge e incident to w is a convex edge. But for any complete vertex v of $\mathcal{C}(F_\eta)$ we must have $v \in S$ and $F_\eta(v) = \eta(v)$, because otherwise the height of v can be increased without violating the local convexity of the surface or the restrictions in η, in contradiction to the maximality of F_η. □

Let us define a *generic* situation (for height vector η and point set S) as one where, for some $\varepsilon > 0$, the maximal locally convex surface for any ε-perturbation η_ε of η and S_ε of S has the same combinatorial structure as the surface F_η.

If the underlying polygonal region R is a convex polygon, then F_η is the lower convex hull [77] of the spatial point set $\{\binom{v}{\eta(v)} \mid v \in S\}$. If, in addition, the situation is generic then $\mathcal{C}(F_\eta)$ is a regular triangulation [64, 40]. This fact generalizes for arbitrary polygonal regions R as follows.

Theorem 7. *Under generic conditions, the complex $\mathcal{C}(F_\eta)$ is a pre-triangulation that is combinatorial regular.*

Proof. By Lemma 21, $F_\eta(v) = \eta(v)$ holds for each complete vertex v of $\mathcal{C}(F_\eta)$. Together with the genericity assumption on η, this implies that the space of exact lifts of $\mathcal{C}(F_\eta)$ has dimension (at least) its degree. Since the underlying point set S is generic, too, it follows that $\mathcal{C}(F_\eta)$ is combinatorial projective. But η is a height vector for $\mathcal{C}(F_\eta)$ such that all edges of F_η are convex. We conclude that $\mathcal{C}(F_\eta)$ is combinatorial regular. Moreover, $\mathcal{C}(F_\eta)$ is a pre-triangulation, by Corollary 5. □

Corollary 8. *If the domain R of F_η is a polygon, then $\mathcal{C}(F_\eta)$ is a face-honest pseudo-triangulation in the generic case.*

Proof. Combine Theorem 6 and Theorem 7. □

4.9 A general flipping scheme

Flip operations are a common tool for locally modifying polygonal complexes. A general class where flip operations have been defined are the pseudo-triangulations. The repertoire includes Lawson flips [62], exchanging flips [75, 87, 88], and removing flips and their inverses [10, 71]. Still, several interesting classes of complexes, in particular the combinatorial regular complexes whose relevance as optimal surfaces is documented in Section 4.8, cannot be reached and modified with these flip operations. Figure 4.15 serves as an example. There exist (generic) height vector(s) for the point set of the triangulation (left) such that the polygonal complex of its maximal locally convex surface is the shown pre-triangulation (right). Observe, that this combinatorial regular complex is not a pseudo-triangulation. Thus the so far defined flips on pseudo-triangulations do not suffice to reach it with a flip sequence starting at the triangulation.

[4]Equivalently, $\varphi(\ell)$ defines a convex function, for every line segment $\ell \subset R$ that does not cross the boundary of R. This definition of locally convexity is used in [10, 25].

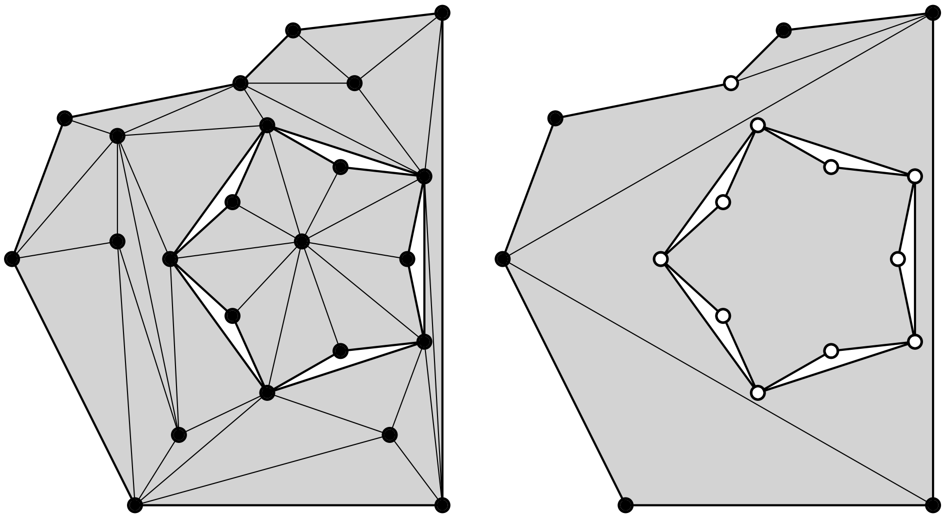

Figure 4.15: There exists a height vector for the depicted point set such that the pre-triangulation at the right hand side is the polygonal complex of the maximal locally convex surface of the triangulation at the left hand side.

Let $\mathcal{C}$ be a face-reducible complex. Below we define a flip operation that applies to any internal edge e of $\mathcal{C}$, in a way such that face-reducibility is retained in the resulting complex. The flip operation is based on locally convex surfaces and on maximal face sets of degree 3. All the flips known for triangulations and pseudo-triangulations arise as special cases of this operation.

Operation FLIP(e). *Let f and g be the two faces of $\mathcal{C}$ incident to e. Consider the subcomplex $\mathcal{C}_e = M_f \cup M_g$ of $\mathcal{C}$. Choose a height for each complete vertex of $\mathcal{C}_e$ such that, if $M_f \neq M_g$, the edge $\varphi|_e$ is reflex in the surface φ for $\mathcal{C}_e$. Let the vector η contain these heights and, in addition, the entry ∞ for each incomplete vertex of $\mathcal{C}_e$. Replace $\mathcal{C}_e$ by $\mathcal{C}(F_\eta)$.*

FLIP(e) deletes a subset E of edges of $\mathcal{C}$ and creates a subset E' of edges disjoint from E. Note that $E \neq \emptyset$ because FLIP(e) always deletes the edge e. That is, each internal edge of $\mathcal{C}$ is flippable with this operation. We might have $E' = \emptyset$ if the flip is removing. The cardinalities of E and E' may be large, though. We offer a detailed discussion of the flip types covered by the operation FLIP in Section 4.10.

Lemma 22. *FLIP(e) constructs a unique and combinatorial regular subcomplex $\mathcal{C}(F_\eta)$.*

Proof. If $M_f = M_g$ then $M_f \cup M_g$ is of degree 3. For any entries in η for these 3 complete vertices, we obviously have $\mathcal{C}(F_\eta)$ unchanged. If $M_f \neq M_g$ then $M_f \cup M_g$ is of degree 4, by Lemma 17. The corresponding 4 entries in η were chosen such that e yields a reflex edge in the surface for $M_f \cup M_g$. This implies that the respective 4 surface vertices are not coplanar. So $\mathcal{C}(F_\eta)$ remains unchanged when η is ε-perturbed. Moreover, even a large change of η does not alter $\mathcal{C}(F_\eta)$ as long as the surface image of e is reflex, because the (three-dimensional) order type of the 4 surface vertices stays the same. This shows that FLIP(e) constructs a unique surface. In particular, η is generic. Moreover, an ε-perturbation of the vertices in $M_f \cup M_g$ will not change the surface combinatorially. Thus generic conditions are given, and $\mathcal{C}(F_\eta)$ is combinatorial regular, by Theorem 7. $\square$

Theorem 8. *The class of face-reducible complexes is closed under the operation FLIP. Pre-triangulations are the smallest class having this property, among the classes listed in Figure 4.3 on page 34.*

Proof. Let $\mathcal{C}$ be some face-reducible complex in the polygonal region $R = U(\mathcal{C})$. Let e be an internal edge of $\mathcal{C}$, and denote with $\mathcal{C}'$ the complex that is obtained from $\mathcal{C}$ by applying FLIP(e).

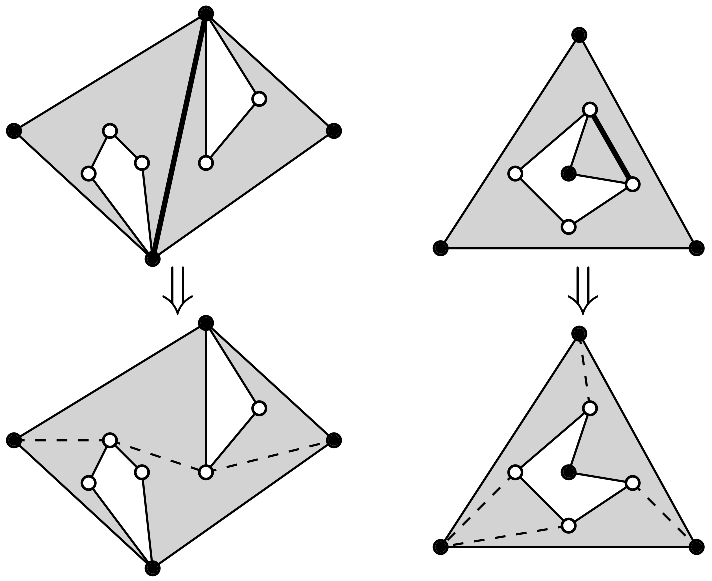

Figure 4.16: Operation FLIP: Two exchanging flips.

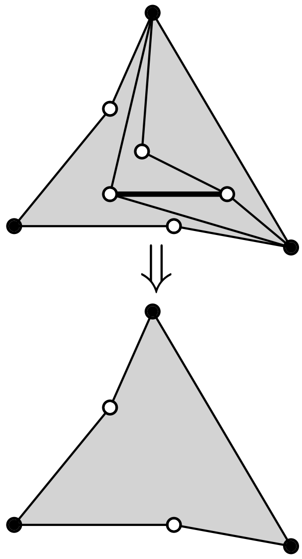

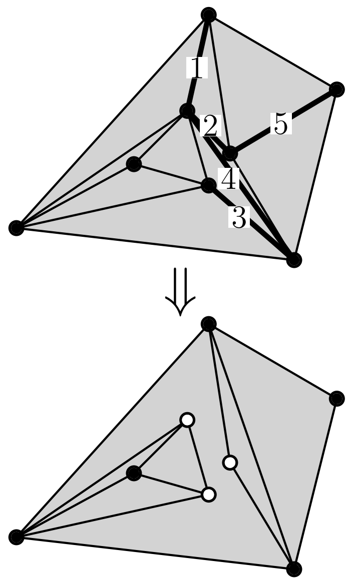

Figure 4.17: Left: Removing flip: Applying FLIP to the bold edge results in the removal of all internal edges and vertices. Right: Many flips: The numbers in the upper example denote a possible flip sequence to reach the lower complex. Flips 1, 2, 3, and 4 are edge removing flips. Flip 5 is edge exchanging.

The subcomplex $\mathcal{C}(F_\eta)$ of $\mathcal{C}'$ constructed by FLIP(e) is combinatorial regular, by Lemma 22. Thus $\mathcal{C}(F_\eta)$ is a pre-triangulation in $Q = U(M_f \cup M_g)$, for f and g being the faces incident to e. In the complement $R \setminus Q$, the complex $\mathcal{C}'$ coincides with $\mathcal{C}$. Moreover, by nature of FLIP(e), the M-skeletons of $\mathcal{C}'$ and $\mathcal{C}$ are identical in $R \setminus Q$. Hence, the part of $\mathcal{C}'$ in $R \setminus Q$ is a face-reducible complex.

If, in addition, $\mathcal{C}$ is a pre-triangulation then there exists a partition of $\mathcal{C}$ into face sets B_i such that $U(B_i)$ is a pre-triangle and no edge of $\mathcal{C}$ is internal to $U(B_i)$, for each set B_i. So B_i is of degree 3, which implies $B_i \subseteq M_b$ if $b \in B_i$ holds for a face b external to Q. This implies that B_i is not affected by FLIP(e) in this case. Consequently, the part of $\mathcal{C}'$ in $R \setminus Q$ is a pre-triangulation.

But the concatenation of two pre-triangulations (or of two face-reducible complexes) is a complex of the same type, as is easily seen from the definition of such complexes.

To complete the proof, we observe that there exist Delaunay triangulations constrained by R that can be flipped to a pre-triangulation which is neither combinatorial projective nor a pseudo-triangulation. Figure 4.17 (right) exemplifies this fact. In view of the containment relations in Figure 4.3 on page 34, this shows that any class closed under FLIP has to include the class of pre-triangulations. □

Note that Theorem 8 does not imply that all face-reducible complexes (or pre-triangulations) within a given region R can be reached by the operation FLIP, if one starts with a fixed one.

4.10 Instances of FLIP

Let us discuss the effect of the operation FLIP in different scenarios. First of all, the interested reader may convince herself that FLIP is able to simulate the exchanging flips in Figure 4.5 and the removing flips in Figure 4.6; see Subsection 4.1.2, page 34 and following. In the figures for the present section (Figures 4.16, 4.17, and 4.18) the edge e to be flipped is shown in bold, and only the region $U(M_f \cup M_g)$ that gets restructured is displayed. Here f and g are the two faces incident to e. Completeness of vertices (indicated by black dots) is with respect to the set $M_f \cup M_g$, and is not necessarily the same in the entire complex. Edges (newly) created by the respective flip are drawn as dashed lines. By Lemma 17, the degree of the face set $M_f \cup M_g$ is at most 4. Therefore, the region $U(M_f \cup M_g)$ (let us call it Q in the discussion below) has either 4 or 3 corners.

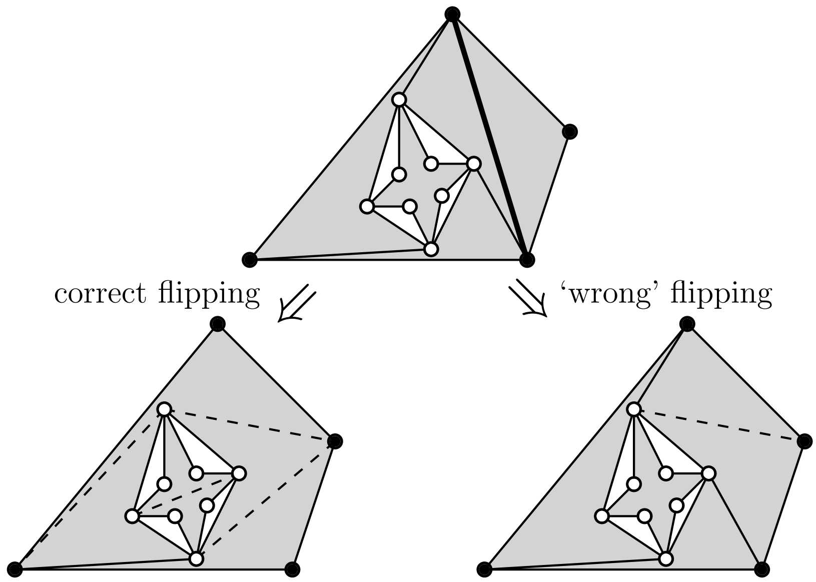

Figure 4.18: FLIP and maximal face sets of degree 3. The edge to be flipped is shown in bold, newly created edges are shown dashed. Left: Correct flip. Right: 'Wrong' flip. The resulting complex is not face-reducible.

The case of 4 corners for Q leads to a generalization of the exchanging flip for pre-triangles; see Figure 4.16. We have $M_f \neq M_g$ and the degree of $M_f \cup M_g$ is 4. As in a usual exchanging flip, the completeness status of all vertices is preserved (apart from possible vertices internal to Q that are necessarily incomplete and are removed by the flip along with their incident edges). More than one edge may be created, as is shown in the left example of Figure 4.16 as well as in the right one. Consequently, FLIP is not a symmetric operation, in general: Flipping one edge (or all the edges in any order) that just have been created need not give the initial complex.

The case of 3 corners for Q constitutes a generalization of the removing flip; see Figure 4.17 (left). All edges and (incomplete) vertices internal to Q get removed. No new edges are created and a single pre-triangle, Q, remains. The degree of $M_f \cup M_g$ is either 3 as in our example, or is 4 as in the flips 1, 2, 3, and 4 in Figure 4.17 (right), and in the examples shown in Figure 4.6 (left and middle) on page 36. This depends on whether $M_f = M_g$ or not. In the negative case, a noncorner of Q which is an endpoint of the flipped edge e changes its status from complete to incomplete.

A natural question is why FLIP(e) is based on maximal face sets of degree 3 rather than on single faces of $\mathcal{C}$. Figure 4.18 provides an answer. The flip at the left is in accordance with the definition of FLIP and retains face-reducibility. In the transformation on the right side, which is solely based on the two faces incident to the flipped edge, face-reducibility is lost, and with it, the applicability of the surface theorem (Theorem 4) and its advantages, which will be exploited next.

4.11 Delaunay minimum complex

By the well-known paraboloid lifting function $\eta(x) = |x|^2$, Delaunay triangulations [21, 43] in a convex region R correspond to maximal locally convex surfaces generated by the vertex heights $\eta(x)$. (In this easy case, such surfaces are just lower convex hulls of the lifted vertices.) This correspondence is extended to non-convex polygons R in [10] and leads to Delaunay triangulations constrained by (the edges of) R. We define below a unique complex of *smallest* combinatorial size in an *arbitrary* polygonal region R that still shows the desired 'Delaunay properties': an analog of local Delaunayhood, and the reachability by improving flip operations.

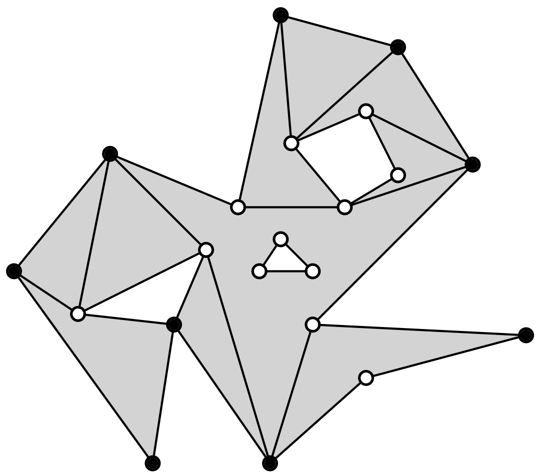

Figure 4.19: Delaunay minimum complex.

Let η be the height vector for the vertices of R such that $\eta(c) = |c|^2$ for each corner c of R, and $\eta(v) = \infty$ for each noncorner v of R. Consider the complex $\mathcal{D}_R = \mathcal{C}(F_\eta)$. We will assume generic conditions for R henceforth, in the sense that ε-perturbations of its vertices (and the changes of vertex heights caused by them) leave the maximal locally convex surfaces above combinatorially unchanged. Then the height vector η is generic, too, and by Theorem 7 the complex $\mathcal{D}_R$ is a pre-triangulation that is combinatorial regular. Moreover, $\mathcal{D}_R$ is a minimum complex in R because, by Lemma 21, its only vertices are those of R and each noncorner of R is incomplete in $\mathcal{D}_R$. We term $\mathcal{D}_R$ the *Delaunay minimum complex* of R. Figure 4.19 gives an example.

Observe that $\mathcal{D}_R$ differs from the constrained Delaunay triangulation [66] of R, unless R is a convex polygon and $\mathcal{D}_R$ is the (classical) Delaunay triangulation of the vertices of R. In case R is a polygon, $\mathcal{D}_R$ is a pointed pseudo-triangulation which is face-honest, by Corollary 8. The latter structure has been also considered in [80] and is a constrained regular pseudo-triangulation as defined in [10].

The class of all possible face-reducible complexes in a given region R is connected under the operation FLIP in the following sense.

Theorem 9. *Let a polygonal region R be given. Each face-reducible complex in R can be transformed to the Delaunay minimum complex $\mathcal{D}_R$ by using finitely many operations FLIP.*

Proof. Let $\mathcal{C}$ be a face-reducible complex in R. We construct a finite sequence of flips that transforms $\mathcal{C}$ into $\mathcal{D}_R$. Let η be a height vector for the vertices of $\mathcal{C}$, with $\eta(c) = |c|^2$ for each corner c of R, and with heights for the remaining vertices of $\mathcal{C}$ being sufficiently large such that $\mathcal{C}(F_\eta) = \mathcal{D}_R$.

Apply η to the complete vertices of $\mathcal{C}$. This defines a unique surface φ on R by Theorem 4. If φ contains reflex edges then choose such an edge $\varphi|_e$, transform $\mathcal{C}$ to a new complex $\mathcal{C}'$ by performing FLIP(e) for the edge e of $\mathcal{C}$, and put $\mathcal{C} = \mathcal{C}'$ and repeat.

The complexes $\mathcal{C}_i$ created above are face-reducible, by Theorem 8. So the procedure is well-defined. We first prove that the created surfaces φ_i are pairwise different[5]. Let FLIP(e) transform $\mathcal{C}_i$ to $\mathcal{C}_{i+1}$. Denote by Q the region restructured by FLIP(e), and let $\mathcal{S}$ and $\mathcal{S}'$, respectively, be the subcomplex of $\mathcal{C}_i$ and $\mathcal{C}_{i+1}$ in Q. By adding all internal edges of $\mathcal{S}'$ to $\mathcal{C}_i$ (and splitting faces of $\mathcal{C}_i$ accordingly) we obtain a polygonal complex A in R. Extend η to a height vector η_A for the complete vertices of A such that the surface for A and η_A is φ_i. (This is possible although A is not face-reducible, in general.) Similarly, extend η to η'_A such that the surface for A and η'_A is φ_{i+1}. Then $\eta'_A \leq \eta_A$ holds (element-wise, with strict inequality for some elements), because e is the only internal edge for $\mathcal{S}$ and $\varphi_i|_e$ is reflex, and all internal edges e' of $\mathcal{S}'$ have a convex image $\varphi_{i+1}|_{e'}$. But surface heights can be shown to strictly decrease with the height vector (see [10], Lemma 5.2). This implies that the surfaces φ_i are pairwise different.

In fact, the complexes $\mathcal{C}_i$ are pairwise different as well: If $\mathcal{C}_i$ and $\mathcal{C}_j$ have different sets of complete vertices then clearly $\mathcal{C}_i \neq \mathcal{C}_j$. Otherwise, by the uniqueness of φ_i for fixed η, $\varphi_i \neq \varphi_j$ implies $\mathcal{C}_i \neq \mathcal{C}_j$. Observe next that no new vertices are created by the operation FLIP. But the number of polygonal complexes with vertices from a fixed set of n points is finite (in fact, exponential in n). We conclude that the procedure above terminates with a surface φ^* where all edges are convex.

In the corresponding complex $\mathcal{C}^*$, the only vertices that are complete are the corners of R: Any other complete vertex v of $\mathcal{C}^*$ would be incident to edges whose images are reflex in φ^*, by the choice of $\eta(v)$. This implies $\varphi^* = F_\eta$ by the uniqueness of F_η. Observe that $\mathcal{C}(F_\eta) = \mathcal{D}_R$ is the M-skeleton of $\mathcal{C}^*$. If $\mathcal{C}^* \neq \mathcal{D}_R$ then $\mathcal{C}^*$ contains internal edges (and possibly, internal and incomplete vertices) that have no counterparts in F_η. As long as such an edge e exists, we apply FLIP(e). For the two faces f and g incident to e, we have $M_f = M_g$. Therefore, FLIP(e) removes all edges internal to $U(M_f)$ (in particular, the edge e) and creates no edges. The number of such flips is bounded by the number of edges of $\mathcal{C}^*$. The surface φ^* remains unchanged, and the complex $\mathcal{D}_R$ is the result. $\square$

Theorem 9 is rather general. Flips can be applied in any order to edges with reflex images, and we even can drop the requirement that such edges have to be flipped before edges without surface counterparts. Moreover, when starting with a pre-triangulation this class is never left; see Theorem 8. We remark that Theorem 9 is a generalization of the optimality theorem for pseudo-triangulations in [10].

As a special case, let R be a polygon, and let PT be some pointed pseudo-triangulation in R without internal vertices. Then PT is face-honest by Corollary 7. Let e be an edge of PT which is incident to two pseudo-triangles ∇ and ∇'. We have $M_\nabla = \nabla \neq \nabla' = M_{\nabla'}$ and the region affected by FLIP(e) is $Q = \nabla \cup \nabla'$. Moreover, Q has exactly 4 corners by Lemma 17. So FLIP(e) is an exchanging flip; see Section 4.10. Let η be the height vector used above to define $D_R = \mathcal{C}(F_\eta)$. We may call FLIP($e$) a *Delaunay flip* if the edge $\varphi|_e$ in the surface φ for PT and η is reflex. As FLIP(e) transforms PT to a pointed pseudo-triangulation of R, and D_R is a complex of this type, we obtain:

[5]Thanks go to Paco Santos for pointing out, in the context of pseudo-triangulations, the following elegant argument.

Corollary 9. *Any pointed pseudo-triangulation of a polygon R can be flipped to D_R by repeated application of Delaunay flips and without leaving the class of pointed pseudo-triangulations.*

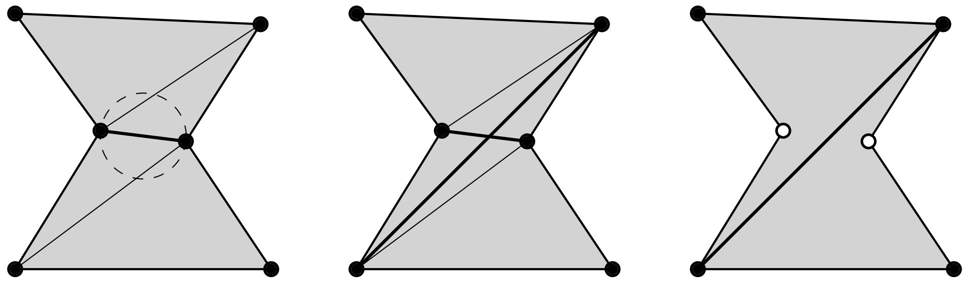

Figure 4.20: The set of edges of the Delaunay minimum complex (right) of the depicted polygon is not a subset of the edge set of its constrained Delaunay triangulation (left).

If Delaunay flips are performed in a well-chosen order then $O(n^2)$ flips are sufficient if R has n vertices, as can be seen by adapting a result in [10]. Note finally that the edges of D_R do not form a subset of the edges of the constrained Delaunay triangulation [66] of R, in general. See Figure 4.20 for an example. The left complex is the Delaunay triangulation constrained by the polygon. The bold edge exists, because there exists a circle (dashed) through its points that is empty of vertices, and the edge is not crossed by any constraint. The pseudo-triangulation to the right depicts the Delaunay minimum complex of the polygon. In this simple example the bold edge is the only internal edge of the complex and it intersects the bold edge of the constrained Delaunay triangulation. (The graph in the middle shows both complexes combined.)

4.12 Face-honest flipping

Throughout this section, we restrict attention to polygonal complexes whose underlying region R is a polygon. We show that any triangulation T in R can be flipped to the complex $\mathcal{C}(F_\eta)$, for any given (generic) height vector η for the vertices of T, in a way such that each intermediate complex is a pseudo-triangulation. (We allow T to contain internal vertices.) In fact, all obtained complexes are face-honest and thus are combinatorial projective, except in cases where an internal vertex has to be removed in the subsequent flip and face-honesty is impossible to keep. This guarantees that only standard flips for pseudo-triangulations are performed. We start with a technical lemma.

Lemma 23. *Let PT be a pseudo-triangulation in R, and let v be an internal and complete vertex of PT with at least 4 incident edges. Then at least one such edge can be flipped by exchange such that no edge incident to v is created.*

Proof. Let Q be the union of the faces of PT incident to v, and let $\mathcal{C}$ be the pseudo-triangulation in Q defined by these faces. For an edge vx of $\mathcal{C}$, define $\alpha(vx)$ as the internal angle at v of the union of the two faces of $\mathcal{C}$ adjacent in vx. Imagine that removing edge flips are applied to all edges vu where u is a noncorner of Q that is complete in $\mathcal{C}$. Each such flip makes a unique vertex incomplete. This vertex is different from v (and thus is u) because $\alpha(vu) < \pi$ holds by our assumption on u. So v stays complete in the obtained complex, and therefore at least 3 edges of $\mathcal{C}$ are still incident to v. Such an edge vw can be flipped in $\mathcal{C}$ by exchange, and without creating an edge incident to v, if $\alpha(vw) < \pi$. But this is guaranteed for at least one such edge, because we assumed v to be incident to at least 4 edges of $\mathcal{C}$. □

Theorem 10. *Let R be a polygon and let T be a triangulation in R. Moreover, let η be a height vector for the vertices of T, and assume generic conditions for the complex $\mathcal{C}(F_\eta)$. Then T can be transformed by FLIP to $\mathcal{C}(F_\eta)$, in a way such that all intermediate complexes are face-honest pseudo-triangulations, except immediately before a vertex is removed.*

Proof. Denote with PT the current complex obtained from flipping; initially, $PT = T$. Recall that all vertices of T are complete. We construct a desired sequence of flips that terminates with $PT = \mathcal{C}(F_\eta)$. Let φ denote the unique surface that results from PT when η is applied to all complete vertices. We say that an internal vertex v of PT fulfills the *hull condition* if $\varphi|_v$ lies strictly below the lower convex hull of its neighbored vertices in φ.

Step (1): As long as there exists an internal vertex v of PT that violates the hull condition, we do the following. Apply exchanging flips to the edges incident to v until only 3 such edges remain. This is possible by Lemma 23. After each flip all the internal vertices of PT are complete, because this was true before the flip, and the flip was exchanging. Two more applications of FLIP to v's remaining edges first make v incomplete (and PT temporary non-projective and hence non face-honest), and then remove v, leaving all internal vertices complete again.

Step (2): While there exists an edge e in PT such that $\varphi|_e$ is reflex, do the following: Apply FLIP(e). If e is exchanged then no vertex alters its completeness status. If e gets removed then an endpoint of e which is a vertex of R becomes incomplete: Each internal vertex fulfills the hull condition before the flip, and $\varphi|_e$ was reflex. So all internal vertices stay complete after FLIP(e). If some internal vertices now violate the hull condition then repeat from Step (1). (Note that PT need not be a triangulation any more.)

The total number of flips performed in Step (1) is clearly $O(n^2)$, if T contains n vertices. In Step (2) only edges with a reflex image are flipped. The number of these flips is finite by the arguments in the proof of Theorem 9. In each created complex all internal vertices are complete, except immediately before a vertex is removed. So, with these exceptions, each such complex is a face-honest pseudo-triangulation, by Lemma 20.

In particular the final complex, PT^*, is face-honest and therefore combinatorial projective. The surface φ^* for PT^* contains no reflex edge. This implies that, for each edge e of PT^*, its image $\varphi^*|_e$ is a convex edge. As only vertices that violate the hull condition are removed, and such vertices cannot belong to $\mathcal{C}(F_\eta)$, we conclude $PT^* = \mathcal{C}(F_\eta)$. $\square$

We conjecture that only quadratically many flips are performed in Step (2). This would yield a flip sequence of total length $O(n^2)$. Theorem 10 implies that the complex $\mathcal{C}(F_\eta)$ for polygons is a pseudo-triangulation. A non-algorithmic proof for this fact has been given in Section 4.8 (Corollary 8).

A result related to Theorem 10 has been proved in [52]. Namely, any triangulation T can be flipped to the complex $\mathcal{C}(F_\eta)$ such that no edges with convex images in the corresponding surface are flipped, and such that all intermediate complexes are combinatorial projective (with the above exceptions before vertex removal).

4.13 Complex classes revisited

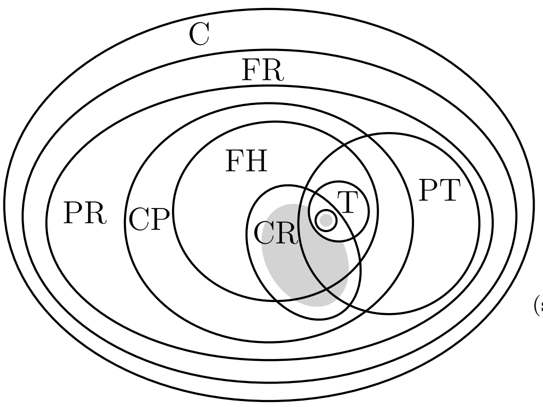

C . . . polygonal complexes (page 38)
FR . . . face-reducible complexes (page 41)
PR . . . pre-triangulations (page 39)
CP . . . combinatorial projective complexes (page 44)
FH . . . face-honest complexes (page 46)
CR . . . combinatorial regular complexes (page 46)
PT . . . pseudo-triangulations
T . . . triangulations
○ . . . Delaunay triangulations constrained by R
(shaded) . . . area for Delaunay minimum complex of R
. . . (for Delaunay minimum complex see page 41)

Figure 4.21: Hierarchy of polygonal complexes in a given region R, revisited.

At the end of this chapter we want to recapitulate the various (new) complex classes. This is meant to give the reader a better understanding of the differences between the complex classes that have been introduced and discussed in the previous sections. For this purpose, we will consider the examples already shown in Figure 4.4 on page 35 in more detail.

For the convenience of the reader we will present all the figures relevant for this section, again. We start with the containment relation diagram (Figure 4.21) of the considered polygonal complexes, already shown in Section 4.1 on page 34. Recall that the picture is non-redundant in the sense that subclasses are proper and class overlaps are nonempty. In the following we will present and explain one example per cell of the diagram.

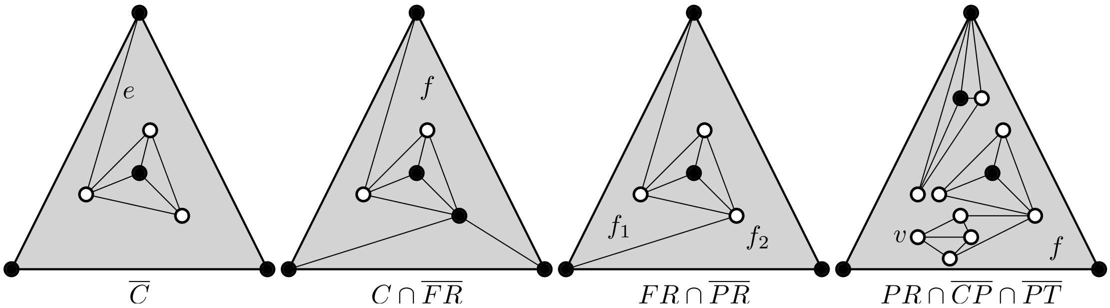

Figure 4.22: The first row of examples from Figure 4.4 on page 35.

We will consider the different class overlaps from the "largest" class to the "smallest" one. We begin with Figure 4.22 from left to right. The leftmost polygonal partition is not a polygonal complex by our definition (see page 38). The interior edge e is adjacent to only one face (at both sides). The second example is a polygonal complex, but not a face-reducible (see page 41) one. There exists no subset of faces of degree 3 containing face f. The third complex is face-reducible but not a pre-triangulation (see page 39). There exists a subset of faces $\{f_1, f_2\}$ with degree 3 containing f_2, but the faces of this subset share edges. The last example shows a pre-triangulation that is neither combinatorial projective (see page 44), because vertex v with all its incident edges will disappear in each lifting, nor a pseudo-triangulation, because face f is not a pseudo-triangle.

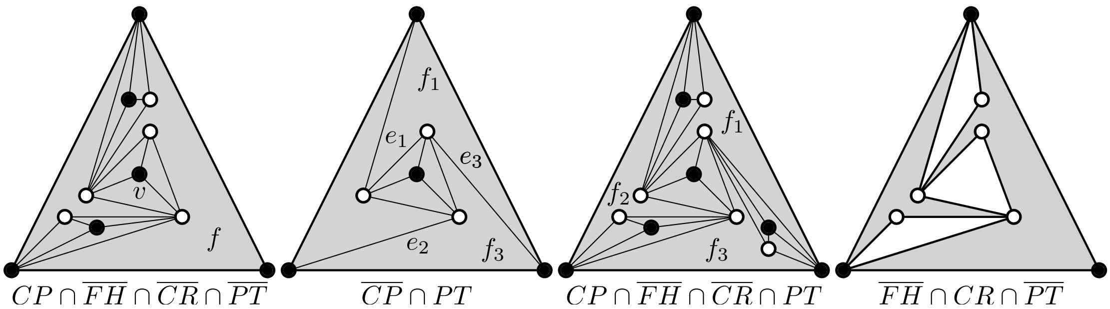

Figure 4.23: The second row of examples from Figure 4.4 on page 35.

Next we consider Figure 4.23, again from left to right. The combinatorial projective complex at the far left is neither face-honest (see page 46), nor combinatorial regular (see page 46), nor a pseudo-triangulation. The face f is not a pre-triangle (thus also no pseudo-triangle) and needs another face to form a pre-triangle (thus $M_f \neq \{f\}$ and the complex is not face-honest). If the edges incident to the complete vertex v are convex in a surface, then the boundary edges of the wheel (see page 11) around v are not convex, and vice versa. Thus the shown complex is also not combinatorial regular. The next example shows a pseudo-triangulation that is not combinatorial projective, as the three edges e_1, e_2, e_3 will not show up in any surface (the three faces f_1, f_2, f_3 are forced to be coplanar in all liftings). Note that $\{f_1, f_2, f_3\}$ is a set of faces with degree 3 that shares edges. But this is not a contradiction to the definition of pre-triangulations, as these three faces are also a "subset of faces of degree 3" for themselves. Next is a combinatorial projective pseudo-triangulation that is neither face-honest (e.g. $\{f_1, f_2, f_3\} = M_{f_1} \neq \{f_1\}$), nor combinatorial regular (each of the 4 wheels of degree 3 witnesses the combinatorial non-regularity of the complex). Knowing the arguments from the examples before, it is easy to see that the rightmost complex is a combinatorial regular pre-triangulation, which is neither face-honest nor a pseudo-triangulation. The two faces are coplanar in every surface and one of the faces is not a pseudo-triangle. The example is combinatorial regular, because each edge is a boundary edge of the underlying polygonal domain (the white triangles are holes in the domain).

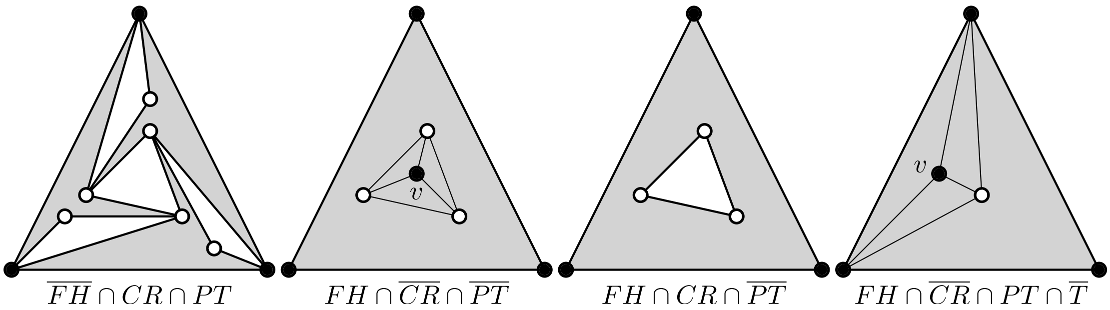

Figure 4.24: The third row of examples from Figure 4.4 on page 35.

Figure 4.24 presents the third set of examples. The leftmost complex is a combinatorial regular (all edges are boundary edges of the underlying polygonal domain) pseudo-triangulation, which is not face-honest (the three pseudo-triangles are coplanar in every surface). The next picture shows a face-honest pre-triangulation, which is neither combinatorial regular nor a pseudo-triangulation. Each of the 4 faces is its own maximal set of degree 3. But one face is not a pseudo-triangle and the wheel of degree 3 (around the vertex v) witnesses the combinatorial non-regularity. The next, very similar example is a pre-triangulation which is face-honest (only one face) and also combinatorial regular (again, all edges are boundary edges of the underlying polygonal domain), but not a pseudo-triangulation (the single face of the complex is not a pseudo-triangle). The last example of this set is a face-honest pseudo-triangulation that is not combinatorial regular (witness: the wheel of degree three around the vertex v) and also not a triangulation (one face is not a triangle).

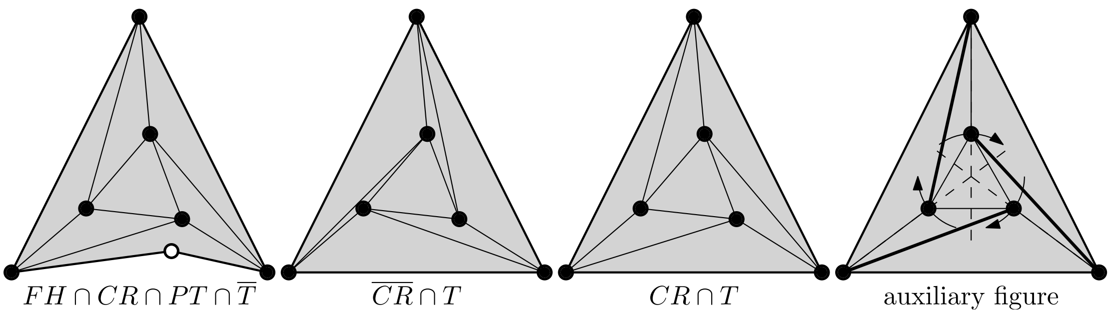

Figure 4.25: The forth row of examples from Figure 4.4 on page 35 except for the rightmost picture, which is an auxiliary drawing for explaining of the combinatorial regularity situation.

The last set of examples (first three leftmost graphs in Figure 4.25) shows three very similar complexes. The leftmost one is a face-honest pseudo-triangulation, the other two are triangulations (which are always face-honest). The first and the third complex (from the left) are combinatorial regular, the second one is not. This can be easily seen by considering the rightmost drawing in Figure 4.25. The vertices of this complex are in degenerate position, in the sense that the three dashed lines intersect in a point. Thus the three bold edges vanish in all possible surfaces. Assume a lifting that makes all edges, except the bold ones, convex. Perturbing the three interior points in the depicted direction (which puts the point set in accordance to the three examples) makes the bold edges visible in the surface, and convex. From the resulting complex, it is easy to see that the first and third complex are combinatorial regular. For the second example, consider the bold edges flipped. With the perturbation these flipped edges become reflex (non-convex). To make them convex, other edges have to become non-convex. Thus this complex is not combinatorial regular.

Concerning the maybe strange appearing markings for Delaunay minimum complexes in Figure 4.21: If the Delaunay minimum complex (see page 51) of a polygonal region R is a triangulation, then it is always a Delaunay triangulation constrained by R. If the Delaunay minimum complex of R is not a triangulation, it still has to be a combinatorial regular complex. But of course, not all combinatorial regular complexes on R are the Delaunay minimum complex of R.

4.14 Chapter summary

We have introduced the concept of pre-triangulations, a relaxation of triangulations that goes beyond the frequently used concept of pseudo-triangulations. Less intuitive at first sight, pre-triangulations turned out to be more natural than pseudo-triangulations in questions concerning liftability and flippability of polygonal complexes – even in the most simple case where the underlying region is a convex polygon.

One of the central tools used in our developments is the completeness status of the vertices of a polygonal complex. (In)completeness of a vertex depends on the (non)convexity of its incident internal angles and thus on the order type of the vertices of the complex in the end. This enabled us to characterize complexes which exhibit the 'robust' lifting property we called combinatorial projectivity.

We did not address algorithmic issues in this chapter. An example would be to decide combinatorial projectivity of a given polygonal complex in an efficient way. In view of Theorem 5 this question reduces to identifying the M-skeleton of the complex or detecting its non-existence. As a simple approach we may try to construct – after ε-perturbing the vertices – a surface for the complex in question, using random heights for its complete vertices. The surface construction mainly means resolving a system of n linear equations [10] where n is the number of vertices. If the system has a solution then the edges of the resulting unique surface correspond to the M-skeleton with probability one (see Corollary 4). Otherwise, the M-skeleton does not exist.

Other relevant algorithmic questions are finding short flip sequences within particular complex classes, and implementing the operation FLIP efficiently. These questions are discussed in detail in a separate work [23]. Emphasis is laid on constructing locally convex functions, and Delaunay-type complexes in particular.

Chapter 5

Conclusion

Considering triangulations and relaxations thereof, we investigated three different topics: Optimality with respect to the minimum total edge length of the complex, flipping in a constrained setting, and lifting the complex while also generalizing the underlying domain to gain brilliant new levels of relaxation for triangulations.

Knowing that pointed pseudo-triangulations can have (depending on the point set) significantly fewer edges than triangulations, it is natural to think that the total edge length of a (pointed) pseudo-triangulation would be less than that of a triangulation. Well, that is not the case in general. We presented examples (of course in addition to the trivial case of points in convex position) where the minimum weight pseudo-triangulation is in fact a triangulation. We also showed that the minimum weight pseudo-triangulation can be a complex that is neither a triangulation nor pointed. On the other hand, we were able to show that a minimum weight pseudo-triangulation of a sufficiently large point set, with approximately at least half of its points in the interior, contains non-extremal pointed vertices. This keeps the hope alive that the weight reduction might be somehow linear in the number of interior vertices. But a good guarantee for the weight loss from relaxing triangulations to pseudo-triangulations on point sets with sufficiently many interior points is still an interesting open question.

That the flip graph of triangulations on point sets is connected is a long standing result. More recent work considers subgraphs of the flip graph of triangulations, asking for the connectedness of the flip graph when only considering triangulations with a certain constraint (for instance the restriction to triangulations that contain a perfect matching). In this respect, we consider pseudo-triangulations that respect a (constant) bound k on the maximum vertex degree. For triangulations on point sets in convex position, we proved connectedness of the flip graph if and only if k is bigger than 6. We also gave an upper bound of $O(n^2)$ on the diameter of this flip graph, but strongly believe that this can be improved. We were able to show an improved bound of $O(n \log n)$ for a slightly relaxed version: When flipping from one triangulation to another one, both respecting a degree bound of k, we allow a degree bound of $k+4$ for intermediate triangulations. We are certain that the $+4$ relaxation for vertex degree bounds of intermediate triangulations can be reduced, maybe even to $+0$. But we leave this for future work.

For point sets in general position we proved the flip graph of degree-bounded triangulations to be disconnected for any k ranging from $\Theta(1)$ to $\Theta(\sqrt{n})$. Whether the flip graph of a relaxed version (like the one for getting a better bound on the convex set described above) would be connected, is a still open question.

Concerning degree-bounded pointed pseudo-triangulations in general position, we showed that their flip graph is not connected as long as the degree bound is at most 9. Whether there exists a (higher) constant degree bound such that the flip graph of degree-bounded (pointed) pseudo-triangulations is connected is an interesting open question. Additional future work may consider the diameter of possible connected flip graphs, relaxed intermediate bounds, or degree-bounded pointed pseudo-triangulations of simple polygons.

Further, for pointed pseudo-triangulations, another question arises when bounding their face degree: the number of vertices that is allowed per pseudo-triangle. The question is trivial for point sets in convex position. But for face degree-bounded pointed pseudo-triangulations on point sets in general position, the connectedness of the flip graph is an interesting open problem; especially if the maximal face degree is bound by 4.

Last but absolutely not least we introduced the new concept of pre-triangulations. We showed that they arise naturally during flipping in lifted triangulations (or pseudo-triangulations). Further, having defined pre-triangulations, lifting and flipping complexes of general polygonal domains becomes possible.

The completeness status of a vertex solely depends on the order type of the underlying point set and is crucial for lifting polygonal complexes. It specifies whether the height of a vertex can be freely chosen, or a vertex' height depends on three (unique) other vertices (and their heights). It is also an indicator of whether or not a (sub)set of faces of a polygonal complex will be coplanar in every possible lifting. Thus the completeness status enables us to characterize polygonal complexes which can be lifted in a "robust" way. We named this class combinatorial projective complexes. "Combinatorial" because, although everything looks very geometric, only combinatorial properties, that is, the two-dimensional order type and the adjacency information of the complex itself, are the determining factors. Following suit, we further defined combinatorial regular complexes, which are a "geometrically stable" version of the well known regular complexes. Finally, for the definition of face-reducible complexes, the completeness status once more is an essential ingredient. Face-reducible complexes define the superclass of all (combinatorial) liftable (in a non-degenerate way) polygonal complexes. Loosely speaking, polygonal complexes which are not face-reducible can only be lifted in a degenerate style (for instance, all complete vertices lie in a plane in $\mathbb{E}^3$, or special "non-stable" geometry), or by "breaking" (inserting additional edges) faces of the complex.

We showed that the flips defined for pseudo-triangulations are inadequate for the newly defined complex classes. Thus we introduced a new flipping scheme, which allows transformation of polygonal complexes as long as they are face-reducible. Pre-triangulations are the smallest class of complexes that are closed under this new flip. Unfortunately, this new flip is neither local nor constant (in the sense of the Lawson-flip for triangulations) in general, but on the plus side extremely versatile. It can "simulate" the basic flip types on pseudo-triangulations while still producing "valid" results when leaving the class of pseudo-triangulations.

Bibliography

[1] M. Abellanas, J. Bose, J. García, F. Hurtado, M. Nicolás, and P. Ramos. On properties of higher order delaunay graphs with applications. In *Proc. 20th European Workshop on Computational Geometry*, pages 119–122, 2005.

[2] P.K. Agarwal, J. Basch, L.J. Guibas, J. Hershberger, and L. Zhang. Deformable free space tilings for kinetic collision detection. In *B.R. Donald, K. Lynch, and D. Rus, editors, Algorithmic and Computational Robotics: New Directions (Proceedings of the 5th Workshop on the Algorithmic Foundations of Robotics)*, 83–96, 2001.

[3] O. Aichholzer, W. Aigner, F. Aurenhammer, T. Hackl, B. Jüttler, E. Pilgerstorfer, and M. Rabl. Divide and conquer for Voronoi diagrams revisited. *Computational Geometry: Theory and Applications*, to appear, 2010. (already appeared in *Proc. 25th Annual ACM Symposium on Computational Geometry*, pages 189–197, 2009.)

[4] O. Aichholzer, F. Aurenhammer, S.-W. Cheng, N. Katoh, G. Rote, M. Taschwer, and Y.-F. Xu. Triangulations intersect nicely. *Discrete & Computational Geometry*, 16(4):339–359, 1996.

[5] O. Aichholzer, F. Aurenhammer, and T. Hackl. Pre-triangulations and liftable complexes. *Discrete & Computational Geometry*, 38:701–725, 2007.

[6] O. Aichholzer, F. Aurenhammer, and T. Hackl. Pre-triangulations and liftable complexes. In *Proc. 22nd Annual ACM Symposium on Computational Geometry*, pages 282–291, 2006.

[7] O. Aichholzer, F. Aurenhammer, T. Hackl, and B. Speckmann. On minimum weight pseudo-triangulations. *Computational Geometry: Theory and Applications*, 42(6-7):627–631, 2009.

[8] O. Aichholzer, F. Aurenhammer, T. Hackl, and B. Speckmann. On (pointed) minimum weight pseudo-triangulations. In *Proc. 19th Annual Canadian Conference on Computational Geometry (CCCG)*, pages 209–212, 2007.

[9] O. Aichholzer, F. Aurenhammer, and H. Krasser. Enumerating order types for small point sets with applications. *Order*, 19:265–281, 2002.

[10] O. Aichholzer, F. Aurenhammer, H. Krasser, and P. Brass. Pseudo-triangulations from surfaces and a novel type of edge flip. *SIAM Journal on Computing* 32:1621–1653, 2003.

[11] O. Aichholzer, S. Cabello, R. Fabila-Monroy, D. Flores-Peñaloza, T. Hackl, C. Huemer, F. Hurtado, and D.R. Wood. Edge-removal and non-crossing configurations in geometric graphs. *Discrete Mathematics & Theoretical Computer Science (DMTCS)*, 12(1):75–86, 2010.

[12] O. Aichholzer, T. Hackl, C. Huemer, F. Hurtado, H. Krasser, and B. Vogtenhuber. On the number of plane geometric graphs. *Graphs and Combinatorics*, 23(1):67–84, 2007.

[13] O. Aichholzer, T. Hackl, D. Orden, P. Ramos, G. Rote, A. Schulz, and B. Speckmann. Flip graphs of bounded-degree triangulations. In *Electronic Notes in Discrete Mathematics 34: Proc. European Conference on Combinatorics, Graph Theory and Applications (EuroComb 2009)*, pages 509–513, 2009.

[14] O. Aichholzer, M. Hoffmann, B. Speckmann, and C.D. Tóth. Degree bounds for constrained pseudo-triangulations. In *Proc. 15th Annual Canadian Conference on Computational Geometry (CCCG)*, pages 155–158, 2003.

[15] O. Aichholzer, C. Huemer, and H. Krasser. Triangulations without pointed spanning trees. *Computational Geometry: Theory and Applications*, 40(1):79–83, 2008.

[16] W. Aigner. The Medial Axis of Planar Shapes. Master Thesis, Institute for Theoretical Computer Science, University of Technology, Graz, Austria, 2007.

[17] T. Asano. Aspect-ratio Voronoi diagram and its complexity bounds. *Information Processing Letters*, 105(1):26–31, 2008.

[18] F. Aurenhammer. A criterion for the affine equivalence of cell complexes in R^d and convex polyhedra in R^{d+1}. *Discrete & Computational Geometry*, 2(1):49–64, 1987.

[19] F. Aurenhammer. Linear combinations from power domains. *Geometriae Dedicata*, 28:45–52, 1988.

[20] F. Aurenhammer. Power diagrams: properties, algorithms, and applications. *SIAM Journal on Computing*, 16(1):78–96, 1987.

[21] F. Aurenhammer. Voronoi diagrams – a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23:345–405, 1991.

[22] F. Aurenhammer and H. Edelsbrunner. An optimal algorithm for constructing the weighted Voronoi diagram in the plane. *Pattern Recognition*, 17(2):251–257, 1984.

[23] F. Aurenhammer, T. Hackl, and H. Krasser. Short flip sequences for constructing the Delaunay triangulation. *Manuscript*, 2010.

[24] F. Aurenhammer, F. Hoffmann, and B. Aronov. Minkowski-type theorems and least-squares clustering. *Algorithmica*, 20:61–76, 1998.

[25] F. Aurenhammer and H. Krasser. Pseudo-simplicial complexes from maximal locally convex functions. *Discrete & Computational Geometry*, 35:201–221, 2006.

[26] F. Aurenhammer and Y.-F. Xu. Optimal triangulations. In *P. Pardalos and C. Floudas, editors, Encyclopedia of Optimization (Second Edition)*, Springer, pages 2757–2764, 2008.

[27] S. Bereg. Transforming pseudo-triangulations. *Information Processing Letters*, 90:141–145, 2004.

[28] M.T. de Berg, O. Cheong, M.J. van Kreveld, and M.H. Overmars. Computational Geometry – Algorithms and Applications (Third Edition). Berlin: Springer-Verlag, 2008.

[29] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. In *D.-Z. Du and F. Hwang, editors, Computing in Euclidean Geometry, Lecture Notes Series on Computing 4*, World Scientific, 47–123, 1995.

[30] P. Bose and F. Hurtado. Flips in planar graphs. *Computational Geometry*, 42:60–80, 2009.

[31] A. Brøndsted. An Introduction to Convex Polytopes. Springer Verlag, 1983.

[32] CGAL. Computational Geometry Algorithms Library. `http://www.cgal.org/`.

[33] B. Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete & Computational Geometry*, 10(1):377–409, 1993.

[34] B. Chazelle, H. Edelsbrunner, M. Grigni, L. Guibas, J. Hershberger, M. Sharir, and J. Snoeyink. Ray shooting in polygons using geodesic triangulations. *Algorithmica*, 12(1):54–68, 1994.

[35] L.P. Chew and R.L. Drysdale. Voronoi diagrams based on convex distance functions. In *Proc. 1st Annual ACM Symposium on Computational Geometry*, pages 235–244, 1985.

[36] H. Crapo and W. Whiteley. Plane self stresses and projected polyhedra. *Structural Topology*, 20:55-78, 1993.

[37] Deutsche Bundesbank, Frankfurt. Homepage of Deutsche Bundesbank — Pictures of and information on the 10 DM banknotes, BBK3 series. `http://www.bundesbank.de/bargeld/bargeld_faq_dmbanknotenabbildungen.en.php`.

[38] H. Edelsbrunner. Algorithms in Combinatorial Geometry. *EATCS Monographs on Theoretical Computer Science*, 10, Springer, 1987.

[39] H. Edelsbrunner and R. Seidl. Voronoi diagrams and arrangements. *Discrete & Computational Geometry*, 1(1):25–44, 1986.

[40] H. Edelsbrunner and N.R. Shah. Incremental topological flipping works for regular triangulations. *Algorithmica*, 15:223–241, 1996.

[41] D. Eppstein. Happy endings for flip graphs. In *Proc. 23rd ACM Symposium on Computational Geometry*, pages 92–101, 2007.

[42] S. Fortune. A sweepline algorithm for Voronoi diagrams. *Algorithmica*, 2:153–174, 1987.

[43] S. Fortune. Voronoi diagrams and Delaunay triangulations. In *D.-Z. Du and F. Hwang, editors, Computing in Euclidean Geometry, Lecture Notes Series on Computing 4*, World Scientific, 225–265, 1995.

[44] S. Fortune. Voronoi diagrams and Delaunay triangulations. In *J.E. Goodman and J. O'Rourke, editors, Handbook of Discrete and Computational Geometry, Second Edition*, CRC Press LLC, 513–528, 2004.

[45] M. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[46] J.E. Goodman and J. O'Rourke, editors. Handbook of Discrete and Computational Geometry, Second Edition. CRC Press LLC, Boca Raton, FL; April 2004.

[47] J.E. Goodman and R. Pollack. Multidimensional sorting. *SIAM Journal Computing* 12:484–507, 1983.

[48] M.T. Goodrich, R. Tamassia. Dynamic Ray Shooting and Shortest Paths in Planar Subdivisions via Balanced Geodesic Triangulations. *Journal of Algorithms*, 23:51–73, 1997.

[49] B. Grünbaum. Convex Polytopes. Wiley Interscience, London, 1967.

[50] J. Gudmundsson and C. Levcopoulos. Minimum weight pseudo-triangulations. *Computational Geometry: Theory and Applications*, 38(3): 139–153, 2007.

[51] R. Haas, D. Orden, G. Rote, F. Santos, B. Servatius, H. Servatius, D. Souvaine, I. Streinu, and W. Whiteley. Planar minimally rigid graphs and pseudo-triangulations. *Computational Geometry: Theory and Applications*, 31(1-2):31–61, 2005.

[52] T. Hackl. Manipulation of Pseudo-Triangular Surfaces. Master Thesis, Institute for Theoretical Computer Science, University of Technology, Graz, Austria, 2004.

[53] S. Hanke, T. Ottmann, S. Schuierer. The edge-flipping distance of triangulations. *Journal of Universal Computer Science*, 2:570–579, 1996.

[54] M. Houle, F. Hurtado, M. Noy, and E. Rivera-Campo. Graphs of triangulations and perfect matchings. *Graphs and Combinatorics*, 21(3):325–331, 2005.

[55] F. Hurtado. Computational/combinatorial/discrete geometry: The flip corner. `http://www-ma2.upc.es/~hurtado/flipcorner.html`.

[56] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete & Computational Geometry*, 22(3):333–346, 1999.

[57] L. Kettner, D. Kirkpatrick, A. Mantler, J. Snoeyink, B. Speckmann, and F. Takeuchi. Tight degree bounds for pseudo-triangulations of points. *Computational Geometry: Theory and Applications*, 12(1-2):1–12, 2003.

[58] D. Kirkpatrick, J. Snoeyink, and B. Speckmann. Kinetic collision detection for simple polygons. *International Journal of Computational Geometry and Applications*, 12(1-2):3–27, 2002.

[59] D. Kirkpatrick and B. Speckmann. Kinetic maintenance of context-sensitive hierarchical representations for disjoint simple polygons. In *Proc. 18th Annual ACM Symposium on Computational Geometry*, pages 179–188, 2002.

[60] V. Klee. On the complexity of d-dimensional Voronoi diagrams. *Archiv der Mathematik*, 34:75–80, 1980.

[61] C.L. Lawson. Generation of a triangular grid with applications to contour plotting. Memo 299, Jet Propulsion Laboratory, 1972.

[62] C.L. Lawson. Properties of n-dimensional triangulations. *Computer Aided Geometric Design*, 3:231–246, 1986.

[63] C.L. Lawson. Transforming triangulations. *Discrete Mathematics*, 3(4):365–372, 1972.

[64] C.W. Lee. Subdivisions and triangulations of polytopes. In *J.E. Goodman and J. O'Rourke, editors, Handbook of Discrete and Computational Geometry, Second Edition*, CRC Press LLC, 383–406, 2004.

[65] D.T. Lee. Two-dimensional Voronoi diagrams in the L_p-Metric. *Journal of the ACM*, 27(4):604–618, 1980.

[66] D.T. Lee and A.K. Lin. Generalized Delaunay triangulation for planar graphs. *Discrete & Computational Geometry*, 1:201–217, 1986.

[67] D.T. Lee and C.K. Wong. Voronoi diagrams in L_1 (L_∞) metrics with 2-dimensional storage applications. *SIAM Journal of Computing*, 9:200–211, 1980.

[68] C. Levcopoulos. An $\Omega(\sqrt{n})$ lower bound for the nonoptimality of the greedy triangulation. *Information Processing Letters*, 25(4):247–251, 1987.

[69] C. Levcopoulos and D. Krznaric. Quasi-greedy triangulations approximating the minimum weight triangulation. *Journal of Algorithms*, 27(2):303–338, 1998.

[70] W. Mulzer and G. Rote. Minimum weight triangulation is NP-hard. *Journal of the ACM*, 55(2), Article 11, 29 pages, 2008.

[71] D. Orden and F. Santos. The Polytope of Non-Crossing Graphs on a Planar Point Set. *Discrete & Computational Geometry*, 33(2):275–305, 2005.

[72] M. Pocchiola and G. Vegter. Minimal tangent visibility graphs. *Computational Geometry: Theorie and Applications*, 6:303–314, 1996.

[73] M. Pocchiola and G. Vegter. On polygonal covers. In *B. Chazelle, J.E. Goodman, and R. Pollack, editors, Advances in Discrete and Computational Geometry, volume 223 of Contemporary Mathematics*, AMS, 257–268, 1999.

[74] M. Pocchiola and G. Vegter, Pseudo-triangulations: theory and applications. In *Proc. 12th Annual ACM Symposium on Computational Geometry*, pages 291–300, 1996.

[75] M. Pocchiola and G. Vegter, Topologically sweeping visibility complexes via pseudo-triangulations. *Discrete & Computational Geometry*, 16:419–453, 1996.

[76] M. Pocchiola and G. Vegter, The visibility complex. *International Journal of Computational Geometry and Applications*, 6(3):279–308, 1996.

[77] F.P. Preparata and M.I. Shamos. Computational Geometry: An Introduction. Springer-Verlag, New York, 1985.

[78] G. Rote, F. Santos, and I. Streinu. Expansive motions and the polytope of pointed pseudo-triangulations. In *B. Aronov, S. Basu, J. Pach, and M. Sharir, editors, Discrete & Computational Geometry - The Goodman-Pollack Festschrift volume 25 of Algorithms and Combinatorics*, Springer, 699–736, 2003.

[79] G. Rote, F. Santos, and I. Streinu. Pseudotriangulations — a survey. In *J.E. Goodman, J. Pach, and R. Pollack, editors, Surveys on Discrete and Computational Geometry: Twenty Years Later, volume 453 of Contemporary Mathematics*, AMS, 343–410, 2008.

[80] G. Rote and A. Schulz. A pointed Delaunay pseudo-triangulation of a simple polygon. In *Proc. 21st European Workshop on Computational Geometry*, pages 77–80, 2005.

[81] G. Rote, C. Wang, L. Wang, and Y. Xu. On constrained minimum pseudotriangulations. In *Proc. 9th Symposium on Computing and Combinatorics*, volume 2697 of *LNCS*, pages 445–454, 2003.

[82] F. Santos. A point set whose space of triangulations is disconnected. *Journal of the American Mathematical Society*, 13(3):611–637, 2000.

[83] R. Sibson. Locally equiangular triangulations. *The Computer Journal*, 21(3):243–245, 1973.

[84] R.I. Silveira. Optimization of polyhedral terrains. PhD thesis, Supervisor: Marc van Kreveld, Utrecht University, 2009.

[85] B. Speckmann and C.D. Tóth. Allocating vertex π-guards in simple polygons via pseudo-triangulations. *Discrete & Computational Geometry*, 33:345–364, 2005.

[86] E. Steinitz. Polyeder und Raumeinteilungen. *Enzyklopaedie der Mathematischen Wissenschaften III AB 12*, Leipzig, 1916.

[87] I. Streinu. A combinatorial approach to planar non-colliding robot arm motion planning. In *Proc. 41st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 443–453, 2000.

[88] I. Streinu. Pseudo-triangulations, rigidity and motion planning. *Discrete & Computational Geometry*, 34:587–635, 2005.

[89] Z o t t e r Schokoladen Manufaktur GmbH, Riegersburg, Austria. Homepage of the z o t t e r chocolate factory. `http://www.zotter.at/en/home.html`.

[90] Z o t t e r Schokoladen Manufaktur GmbH, Riegersburg, Austria. Homepage of Mitzi Blue®. `http://www.mitziblue.com/`.
`http://www.zotter.at/en/chocolate-shop/mitzi-blue/detail/v/kuerbiskern-1.html`.

Curriculum Vitae of Thomas Hackl

April 12, 2010

Affiliation and address

Dipl.-Ing. Thomas Hackl
Research Assistant
Institute for Software Technology
Graz University of Technology
Inffeldgasse 16b, II
A-8010 Graz, Austria
Phone: ++43 316 873 5702
Email: thackl@ist.tugraz.at

Personal Information

Born July 4, 1976 in Steyr (Austria)
Austrian nationality

Education

Graz University of Technology, Austria
2005 - present
PhD student, Advisors: Prof. O. Aichholzer, Prof. F. Aurenhammer

Graz University of Technology, Austria
1996 - 2004
Dipl.-Ing. (MSc), Telematics, Advisors: Prof. F. Aurenhammer
Studies of Telematics, Diploma Thesis at the Institute for Theoretical Computer Science (with distinction)

Higher Technical School for Electronics & Information Technology, Steyr, Austria
1990 - 1995
Matura (A-levels) in digital electronics, information technology, measurement-, systems- and control engineering

Professional Experience

Graz University of Technology, Austria
2005 - present
Institute for Software Technology
Research Assistant, supported by the Austrian FWF National Research Network 'Industrial Geometry' S9205-N12.

COMYAN Internet & Intranet Solutions GmbH
2000 - 2003
Research and Development Center, Graz, Austria
Part time software engineer

Other Occupation

Austrian Armed Forces
1995 - 1996
national service (eight months)

Career Related Activities

Program Committee	PC-member of the 23rd European Workshop on Computational Geometry 2007, Graz, Austria.
Organization	Since 2008 collaborator in the Coordination Project NRN S9201-N12. Co-organizer of the annual European Pseudo-triangulation research week in 2008 and 2009. Co-chair of the organizing committee of the 23rd European Workshop on Computational Geometry 2007, Graz, Austria.
Co-supervisor	Of 4 Diploma thesis and over 15 programming projects on algorithms, geometry, and optimization.
Referee	For Computational Geometry: Theory and Applications (Elsevier), IEEE Transactions on Visualization and Computer Graphics, Information Processing Letters (Elsevier), the 23rd European Workshop on Computational Geometry (EuroCG 2007), and the 9th Latin American Theoretical Informatics Symposium (LATIN 2010).
Research stays	Universidad de Alcalá, Alcalá de Henares (Madrid, Spain), Freie Universität Berlin (Germany), TU Eindhoven (Netherlands), Utrecht University (Netherlands).

Research Interests

Data structures and algorithms in general, graph & enumeration algorithms and computational & combinatorial geometry in particular, especially triangulations, pseudo-triangulations and related structures.

Publications overview

9+2+1 Articles in refereed journals. (published + to appear + submitted)
19 Articles in refereed proceedings.
Editor of **1** proceedings.
Diploma Thesis.

The publications marked bold in the following list arose from parts of the doctoral thesis of Thomas Hackl. The list is in reverse chronological order (sorted alphabetically by authors within each year) and categorized in journal appearances, conference papers, editor of proceedings, and theses.

Publications

9+2+1 articles in refereed journals:

submitted

O. Aichholzer, T. Hackl, M. Hoffmann, C. Huemer, F. Santos, B. Speckmann, and B. Vogtenhuber. Maximizing maximal angles for plane straight line graphs. *submitted to journal*, 2010.

to appear

O. Aichholzer, W. Aigner, F. Aurenhammer, T. Hackl, B. Jüttler, E. Pilgerstorfer, and M. Rabl. Divide-and conquer for Voronoi diagrams revisited. *Computational Geometry: Theory and Applications*, page to appear, 2010.

O. Aichholzer, F. Aurenhammer, T. Hackl, B. Jüttler, M. Oberneder, and Z. Sír. Computational and structural advantages of circular boundary representation. *International Journal of Computational Geometry & Applications*, page to appear, 2010. Special Issue.

2010

O. Aichholzer, S. Cabello, R. Fabila-Monroy, D. Flores-Peñaloza, T. Hackl, C. Huemer, F. Hurtado, and D.R. Wood. Edge-removal and non-crossing configurations in geometric graphs. *Discrete Mathematics & Theoretical Computer Science (DMTCS)*, 12(1):75–86, 2010.

O. Aichholzer, T. Hackl, C. Huemer, F. Hurtado, and B. Vogtenhuber. Large bichromatic point sets admit empty monochromatic 4-gons. *SIAM Journal on Discrete Mathematics (SIDMA)*, 23(4):2147–2155, 2010.

2009

O. Aichholzer, W. Aigner, F. Aurenhammer, T. Hackl, B. Jüttler, and M. Rabl. Medial axis computation for planar free-form shapes. *Computer-Aided Design*, 41(5):339–349, 2009. Special issue: Voronoi Diagrams and their Applications.

O. Aichholzer, F. Aurenhammer, T. Hackl, and B. Speckmann. On minimum weight pseudo-triangulations. *Computational Geometry: Theory and Applications*, 42(6-7):627–631, 2009.

O. Aichholzer, R. Fabila-Monroy, D. Flores-Peñaloza, T. Hackl, C. Huemer, and J. Urrutia. Empty monochromatic triangles. *Computational Geometry: Theory and Applications*, 42(9):934–938, 2009.

2008

O. Aichholzer, F. Aurenhammer, P. Gonzalez-Nava, T. Hackl, C. Huemer, F. Hurtado, H. Krasser, S. Ray, and B. Vogtenhuber. Matching edges and faces in polygonal partitions. *Computational Geometry: Theory and Applications*, 39(2):134–141, 2008.

2007

O. Aichholzer, F. Aurenhammer, and T. Hackl. Pre-triangulations and liftable complexes. *Discrete & Computational Geometry*, 38:701–725, 2007.

O. Aichholzer, F. Aurenhammer, T. Hackl, and C. Huemer. Connecting colored point sets. *Discrete Applied Mathematics*, 155(3):271–278, 2007.

O. Aichholzer, T. Hackl, C. Huemer, F. Hurtado, H. Krasser, and B. Vogtenhuber. On the number of plane geometric graphs. *Graphs and Combinatorics (Springer)*, 23(1):67–84, 2007.

19 articles in refereed proceedings:

2009

O. Aichholzer, W. Aigner, F. Aurenhammer, T. Hackl, B. Jüttler, E. Pilgerstorfer, and M. Rabl. Divide-and-conquer for voronoi diagrams revisited. In 25^{th} *Annual ACM Symposium Computational Geometry*, pages 189–197, Aarhus, Denmark, 2009.

O. Aichholzer, W. Aigner, F. Aurenhammer, T. Hackl, B. Jüttler, E. Pilgerstorfer, and M. Rabl. Divide-and-conquer for voronoi diagrams revisited. In *Proc.* 25^{th} *European Workshop on Computational Geometry EuroCG '09*, pages 293–296, Brussels, Belgium, 2009.

O. Aichholzer, F. Aurenhammer, O. Devillers, T. Hackl, M. Teillaud, and B. Vogtenhuber. Lower and upper bounds on the number of empty cylinders and ellipsoids. In *Proc.* 25^{th} *European Workshop on Computational Geometry EuroCG '09*, pages 139–142, Brussels, Belgium, 2009.

O. Aichholzer, R. Fabila-Monroy, D. Flores-Peñaloza, T. Hackl, C. Huemer, J. Urrutia, and B. Vogtenhuber. Modem illumination of monotone polygons. In *Proc.* 25^{th} *European Workshop on Computational Geometry EuroCG '09*, pages 167–170, Brussels, Belgium, 2009.

O. Aichholzer, T. Hackl, M. Hoffmann, A. Pilz, G. Rote, B. Speckmann, and B. Vogtenhuber. Plane graphs with parity constraints. In *Lecture Notes in Computer Science, Proc.* 11^{th} *International Workshop on Algorithms and Data Structures (WADS)*, volume 5664, pages 13–24, Banff, Alberta, Canada, 2009.

O. Aichholzer, T. Hackl, C. Huemer, F. Hurtado, and B. Vogtenhuber. Large bichromatic point sets admit empty monochromatic 4-gons. In *Proc.* 25^{th} *European Workshop on Computational Geometry EuroCG '09*, pages 133–136, Brussels, Belgium, 2009.

O. Aichholzer, T. Hackl, D. Orden, P. Ramos, G. Rote, A. Schulz, and B. Speckmann. Flip graphs of bounded-degree triangulations. In *Electronic Notes in Discrete Mathematics: Proc. European Conference on Combinatorics, Graph Theory and Applications EuroComb 2009*, volume 34, pages 509–513, Bordeaux, France, 2009.

2008

O. Aichholzer, F. Aurenhammer, T. Hackl, B. Kornberger, S. Plantinga, G. Rote, A. Sturm, and G. Vegter. Seed polytopes for incremental approximation. In *Proc.* 24^{th} *European Workshop on Computational Geometry EuroCG '08*, pages 13–16, Nancy, France, 2008.

O. Aichholzer, S. Cabello, R. Fabila-Monroy, D. Flores-Peñaloza, T. Hackl, C. Huemer, F. Hurtado, and D.R. Wood. Edge-removal and non-crossing configurations in geometric graphs. In *Proc.* 24^{th} *European Workshop on Computational Geometry EuroCG '08*, pages 119–122, Nancy, France, 2008.

O. Aichholzer, R. Fabila-Monroy, D. Flores-Peñaloza, T. Hackl, C. Huemer, and J. Urrutia. Empty monochromatic triangles. In *Proc.* 20^{th} *Annual Canadian Conference on Computational Geometry CCCG 2008*, pages 75–78, Montreal, Quebec, Canada, 2008.

2007

O. Aichholzer, F. Aurenhammer, T. Hackl, B. Jüttler, M. Oberneder, and Z. Sír. Computational and structural advantages of circular boundary representation. In *Lecture Notes in Computer Science, Proc.* 10^{th} *International Workshop on Algorithms and Data Structures (WADS)*, volume 4619, pages 374–385, Halifax, Nova Scotia, Canada, 2007.

O. Aichholzer, F. Aurenhammer, T. Hackl, B. Kornberger, M. Peternell, and H. Pottmann. Approximating boundary-triangulated objects with balls. In *Proc.* 23^{rd} *European Workshop on Computational Geometry EuroCG '07*, pages 130–133, Graz, Austria, 2007.

O. Aichholzer, F. Aurenhammer, T. Hackl, and B. Speckmann. On (pointed) minimum weight pseudo-triangulations. In *Proc.* 19^{th} ***Annual Canadian Conference on Computational Geometry CCCG 2007*, pages 209–212, Ottawa, Ontario, Canada, 2007**.

O. Aichholzer, T. Hackl, M. Hoffmann, C. Huemer, F. Santos, B. Speckmann, and B. Vogtenhuber. Maximizing maximal angles for plane straight line graphs. In *Proc.* 23^{rd} *European Workshop on Computational Geometry EuroCG '07*, pages 98–101, Graz, Austria, 2007.

O. Aichholzer, T. Hackl, M. Hoffmann, C. Huemer, A. Por, F. Santos, B. Speckmann, and B. Vogtenhuber. Maximizing maximal angles for plane straight line graphs. In *Lecture Notes in Computer Science, Proc.* 10^{th} *International Workshop on Algorithms and Data Structures (WADS)*, volume 4619, pages 458–469, Halifax, Nova Scotia, Canada, 2007.

2006

O. Aichholzer, F. Aurenhammer, and T. Hackl. Pre-triangulations and liftable complexes. In 22^{nd} ***Annual ACM Symposium Computational Geometry*, pages 282–291, Sedona, Arizona, USA, 2006**.

O. Aichholzer, T. Hackl, C. Huemer, F. Hurtado, H. Krasser, and B. Vogtenhuber. On the number of plane graphs. In *Proc.* 17^{th} *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 504–513, Miami, Florida, USA, 2006.

2005

O. Aichholzer, F. Aurenhammer, P. Gonzalez-Nava, T. Hackl, C. Huemer, F. Hurtado, H. Krasser, S. Ray, and B. Vogtenhuber. Matching edges and faces in polygonal partitions. In *Proc.* 17^{th} *Annual Canadian Conference on Computational Geometry CCCG 2005*, pages 123–126, Windsor, Ontario, Canada, 2005.

O. Aichholzer, T. Hackl, C. Huemer, F. Hurtado, H. Krasser, and B. Vogtenhuber. Bounding the number of plane graphs. In *Proc.* 15^{th} *Annual Fall Workshop on Computational Geometry and Visualization*, pages 31–32, Philadelphia, Pennsylvania, USA, 2005.

Editor of **1** proceedings:

O. Aichholzer and T. Hackl, editors. *Collection of Abstracts of the* 23^{rd} *European Workshop on Computational Geometry 2007*, Graz, Austria, 2007.

Thesis:

T. Hackl. Manipulation of pseudo-triangular surfaces. Master's thesis, IGI-TU Graz, Austria, 2004.