

Manipulation of Pseudo-Triangular Surfaces

Diploma Thesis
at
Graz University of Technology

submitted by
Thomas Hackl

*Institute for Theoretical Computer Science,
Graz University of Technology
A-8010 Graz, Austria*

Supervisor: Dipl.-Ing. Dr.techn. Univ.Doiz. Univ.-Prof. Franz Aurenhammer
Advisor: Dipl.-Ing. Dr.techn. Univ.Doiz. Univ.-Prof. Franz Aurenhammer
Dipl.-Ing. Dr.techn. Oswin Aichholzer

©Copyright 2004 by Thomas Hackl

Abstract

This diploma thesis deals with pseudo-triangular surfaces and flipping therein, as introduced in [3]. They defined a *projectivity* attribute for pseudo-triangulations and introduced a *stability* condition to decide it. Using a program from preliminary work of this thesis, we found a counter-example for concluding from stability to projectivity.

Our aim is to redefine the stability-condition to be able to correctly conclude to projectivity in all cases. Our investigations lead to a proper combinatorial understanding of the projectivity of pseudo-triangulations. Thereby, we find a new class of cell complexes: *punched pseudo-triangulations*, which are a relaxation of pseudo-triangulations. In addition, we prove the existence of finite flipping sequences to the optimal surface that avoid the creation of non-pseudo-triangular cell complexes.

Keywords: Computational geometry, pseudo-triangulation, flip operations, surface realization, locally convex function, punched pseudo-triangulation, 3-reducibility, projectivity, stability

I hereby certify that the work presented in this thesis is my own and that work performed by others is appropriately cited.

Thomas Hackl
Graz, Austria, September 2004

Acknowledgements

My best thanks go to my parents, without their support and understanding this thesis would not have been possible.

I am indebted to my thesis supervisor Prof. Dipl.-Ing. Dr. Franz Aurenhammer, for introducing me to the topic, for his advises, for the long discussions and for reviewing this thesis.

I would also like to express thanks to Dipl.-Ing. Dr. Oswin Aichholzer for his advices and suggestions. I am also grateful for him for reviewing this thesis.

My thank goes to Dipl.-Ing. Dr. Hannes Krasser for his comments and help during the preceding project that provides a graphical visualization of the underlying paper and also this thesis.

Contents

1	Introduction	1
1.1	Overview of results	2
2	Basic Theory	3
2.1	Overview of basic structures	3
2.2	Planar pseudo-triangulations in a nutshell	5
2.3	Planar flipping	7
2.4	Surfaces for pseudo-triangulations	12
3	3-Reducibility	21
3.1	The problem	21
3.2	3-Reducibility	22
3.3	Algorithm to decide 3-reducibility	26
4	Two Implications of 3-Reducibility	29
4.1	Sets with 3 height defining vertices	29
4.2	Characterizing combinatorial projectivity	30
5	Combinatorial Stability	34
5.1	Finding the correct complete vertices	34
5.2	The updated stability property	39
6	Appearance of Hidden Edges	43
6.1	The deformation	46
6.2	Leaving deformations	48
6.3	Bypassing deformations	51
7	Summary	55
7.1	Application of theory	55
7.2	Conclusion	55
7.3	Future Work	56
A	Flip Sequence to Optimum	57
B	Surface Examples with Deformation	61

List of Figures

1	Outer boundary of polygonal regions.	3
2	Edge-exchanging flip types and non-flipable edge.	7
3	Ambiguous geodesics interpretation.	9
4	Vertex-removing flip.	10
5	The difference between the concepts of projectivity and combinatorial projectivity.	14
6	Two convexifying flips.	17
7	A planarizing flip.	17
8	Example where each combinatorial planar edge is incident to some combinatorial planar vertex.	18
9	Local and global optima.	19
10	(Deformed) double adjacencies.	21
11	Global dependency of vertex heights	22
12	Two detailed back-tracing examples	26
13	Combinatorial projective pseudo-triangulation containing combinatorial coplanar facets.	31
14	Stability vs. combinatorial projectivity.	34
15	A complete vertex, v , within a pseudo-triangle, is convex encircled with incompletes.	35
16	With incompletes non-convex encircled complete vertex.	36
17	Two complete vertices that are complete encircling each other.	38
18	Edges of ∇ are not admissible for valid encirculations.	39
19	Vertices of ∇ are admissible for valid encirculations.	39
20	Double adjacencies are ruled out in combinatorial stable \mathcal{PT} 's.	41
21	Reappearance of a hidden edge.	44
22	Example for a simple deformation.	46
23	The three cases of an improving surface flip.	47
24	Example of the 2 phases of a deformation creation.	47
25	Edge-flip on the outer boundary of a maximal punched set.	49
26	The different edge categories for creating a deformation.	51
27	Three edges keep a vertex complete.	52
28	Improving surface flip sequence to optimum, 01/04.	57
29	Improving surface flip sequence to optimum, 02/04.	58
30	Improving surface flip sequence to optimum, 03/04.	59
31	Improving surface flip sequence to optimum, 04/04.	60
32	Example surface with deformation during flipping to optimum, 01/02.	61
33	Example surface with deformation during flipping to optimum, 02/02.	62

1 Introduction

In computational geometry [21] it is a common task to partition a given geometric domain into cell complexes. This is often done using triangular meshes, usually referred to as triangulations. In the last years, a relaxation of triangulations, the so-called pseudo-triangulations, attracted more and more interest because of their large area of application. Pseudo-triangulations apply to ray shooting [9, 12], visibility [19, 18], guarding [20, 23], rigidity [24, 22, 13] and kinetic collision detection [1, 14, 15].

The article [3], the groundwork for this thesis, introduces polyhedral surfaces in 3-D for pseudo-triangulations. The paper proved that each pseudo-triangulation can be "lifted" into a surface and adapted the 2-dimensional edge flips for pseudo-triangulations to surfaces of pseudo-triangulations. Further, the work [3] introduced the projectivity attribute for pseudo-triangulations, because the surfaces of projective pseudo-triangulations always project to cell complexes that again are pseudo-triangulations in the generic case.

To decide projectivity, the stability condition for pseudo-triangulations was introduced in [3]. Excluding special cases based on the geometric placement, the stability allowed to directly conclude to projectivity. They were also able to show the existence of a finite sequence of improving surface flips to the optimal surface from every given starting triangulation.

In the preliminary stages of this thesis a computer program was developed, using the theory provided in [3]. This program was designed to calculate and display surfaces from pseudo-triangulations and to allow the user to apply surface flips by clicking edges of the surface. During the test phase, there sometimes emerged surfaces that were not predicted by the theory, although the program worked correctly with respect to the theory. Though the underlying pseudo-triangulations of the surfaces were stable, the projections of those surfaces were no pseudo-triangular cell complexes any more. Therefore further flipping was not possible, as flips were only defined for pseudo-triangulations, respectively surfaces thereof.

After closer examination we discovered that we had found examples where the stability-condition from [3] does not allow to correctly decide projectivity of the pseudo-triangulation. We searched for other counter-examples for the stability-projectivity relation and tried to classify them. We paid attention to commonalities and attempted to formulate simple conditions for a new and correct stability. But we soon faced the fact that the problem was not as easy as thought at the beginning.

Therefore we restarted at investigating the projectivity of pseudo-triangulations and searching for new attributes and conditions that describe projectivity. We investigated combinatorial coplanarity of facets, what is a necessary requirement for non-projectivity of pseudo-triangulations.

When we successfully corrected the stability-condition and thereby discovered that flipping may lead into non-pseudo-triangular cell complexes, we were concerned that a finite flipping sequence to the optimal surface will not be guaranteed within

the class of pseudo-triangulations. Basically the problem arose, what to do, when flipping into such a cell complex. So we had to analyze the non-pseudo-triangular cell complexes and show that it is always possible to reach the optimal surface with a finite sequence of admissible flips.

1.1 Overview of results

We start by stating some geometric definitions, forming the fundamental theory used for this thesis, in Section 2. This includes basic geometrical structures like the convex hull and triangulations (Section 2.1), pseudo-triangulations and their properties (Section 2.2) and flipping therein (Section 2.3). Further we summarize the results and definitions about surfaces from pseudo-triangulations from [3] and state some additional definitions for later use, in Section 2.4.

In Section 3, we first exemplify the insufficiency of the stated stability-projectivity connection from [3]. We show that the phenomenon is forced by combinatorial planar edges that cannot be removed by trivial flips. We investigate the connection between these edges and combinatorial coplanar subsets of faces and introduce the concept of *3-reducibility*. In Section 4 we use the 3-reducibility as a powerful tool for deciding combinatorial coplanarity and (in an advanced form) also for deciding combinatorial projectivity.

In Section 5, we investigate the different structures that form combinatorial non-projective pseudo-triangulations. Finally we state the definition for the new stability condition, called *combinatorial stability*, and prove that this condition suffices to characterize combinatorial projectivity.

In the last section (Section 6) we prove the existence of a finite sequence of admissible flips to the optimal surface. Therefore we analyze the special case of combinatorial non-projective pseudo-triangulations, where the projection of the surface is a *punched pseudo-triangulation*, where so-called *deformations* arise. We introduce the *hidden edges*, combinatorial planar edges that remain after a trivial flip in the case of deformations. First we reprove the existence of a finite flipping sequence to the optimal surface, by using hidden edges to complete punched pseudo-triangulations to "real" pseudo-triangulations. Further we provide an upper bound on the number of hidden edges. Finally, we prove the existence of a finite sequence of improving surface flips (including the trivial flip) from every triangular surface to the optimum, without ever creating a deformation and therefore never needing hidden edges.

2 Basic Theory

2.1 Overview of basic structures

At first we want to (re)state some fundamental geometric definitions that will be used in this thesis. This should prevent misunderstandings as some definitions may slightly differ from other publications.

We start by defining the *convex hull* of a set, S , of points. There exist a lot of different equivalent definitions for the convex hull. A natural concept that is easy to understand is the so-called rubber band definition: "If S is a finite set of points in the plane, imagine surrounding S by a large, stretched rubber band; when the band is released it will assume the shape of the convex hull of S " [21]. A more mathematical and general definition that can also be found in [21] is the following one. We denote the d -dimensional Euclidean space with \mathbb{E}^d .

Definition 2.1 (Convex Set) *A subset D of \mathbb{E}^d is convex if, for each pair of points (p_1, p_2) in D , the straight line segment $\overline{p_1 p_2}$ is entirely contained in D .*

Definition 2.2 (Convex Hull) *The convex hull, denoted as $\text{conv}(S)$, of a finite set, S , of points is the smallest convex set containing S .*

In this thesis we are basically operating in 2-dimensional space and enter 3-space only for surface construction, what will be discussed later on. Therefore we restrict ourselves to \mathbb{E}^2 for the next definitions. Similar definitions can be found in [21, 8, 7].

Definition 2.3 (Polygonal Region) *A polygonal region, R , is a 2-dimensional subset of the plane with piecewise linear boundary, denoted with $\partial(R)$. The 1-dimensional components of $\partial(R)$ are called edges of R . The endpoints of edges of R (the 0-dimensional components of $\partial(R)$) are called vertices of R . The set of vertices of R is denoted with $\text{vert}(R)$.*




Figure 1: Outer boundary of polygonal regions.

A polygonal region may be disconnected, and also may contain holes. Therefore it is sometimes advantageous to be able to refer only to the outer boundary of the polygonal region.

Definition 2.4 (Outer Boundary) *The outer boundary of a polygonal region, R , consists of the points of $\partial(R)$ which can be connected to infinity with a curve that intersects R in no other point.*

Figure 1 shows different polygonal regions (gray) and their outer boundaries (bold). Vertices that are on the outer boundary are white, others black. In Figure 1(a) each vertex of the polygonal region is on the outer boundary, as we can draw a curve from each point to infinity. In Figure 1(b) the polygonal region is split into two parts, one of them having a hole. The outer boundary of the polygonal region is formed by vertices of both parts. But as it is not possible to draw a curve from the vertices of the hole to infinity without intersecting with the polygonal region, the vertices of the hole do not belong to the outer boundary. The polygonal region in Figure 1(c) consists also of two parts. But now one part is within the hole of the other one. The displayed connection to infinity (dashed curve) is intersecting with the polygonal region and therefore not allowed.

Note that an edge with only one vertex on the outer boundary need not belong to the outer boundary. The area, enclosed by the outer boundary of a polygonal region, is not necessarily the interior of this polygonal region (think of holes). But the interior is always enclosed by the outer boundary.

Definition 2.5 ((Simple) Polygon) *A (simple) polygon, P , is a polygonal region, homeomorphic to a disk.*

Note that the convex hull of a point set, S , is a convex polygon. In particular, $\text{conv}(S)$ is the smallest convex polygon covering S . Therefore the vertices of the convex hull of S are denoted by $\text{vert}(\text{conv}(S))$.

Definition 2.6 (Cell Complex) *Let R be a polygonal region. A cell complex, $\mathcal{C}(R)$, in R is a polygonal partition of R , such that two cells (faces) of $\mathcal{C}(R)$ either do not intersect or their intersection consists of edges or vertices (faces of lower dimension) of $\mathcal{C}(R)$ which both faces have in common.*

To avoid discussion of special cases we assume general position, i.e., no three collinear vertices, of the vertex sets used in this thesis. Further, we consider only simply connected polygonal regions as domains for cell complexes, throughout this thesis. Nevertheless, subsets of cell complexes may form polygonal regions with internal holes. The more general case of domains that may contain holes will be subject of further investigation in the future.

Definition 2.7 (Triangulation) *Let R be a polygonal region. A triangulation, $T(R)$, in R is a cell complex in R with exclusively triangular faces.*

Normally triangulations are defined on finite point sets like: A triangulation, $T(S)$, of a finite set S of points in \mathbb{E}^2 is a maximal planar straight-line graph that

uses exactly the points of S as its vertices. [21]. With Definition 2.7, a triangulation, $T(S)$, of a finite set S of points is simply a triangulation, $T(R)$, in a polygonal region R , where R equals the convex hull of S and $T(R)$ uses exactly the points of S as vertices. The reason why we are using a different definition for triangulations (and later for pseudo-triangulations) is that we need much more general underlying domains than the convex hull, for instance, an arbitrary simple polygon with possible holes in the interior that forms the boundary for a finite set of points.

2.2 Planar pseudo-triangulations in a nutshell

A relaxation of triangulations are the so-called *pseudo-triangulations* that consist of *pseudo-triangles*. In this section we want to give a short introduction to this topic to get the reader acquainted, as pseudo-triangulations are an important tool for flipping surfaces to optimality, as we will see later on.

Pseudo-triangles are structures similar to triangles, i.e.: the convex hull of a pseudo-triangle is a triangle. Unlike triangles, pseudo-triangles may exist having more than three vertices. To define pseudo-triangles we need a few attributes of polygons that will be specified now.

Definition 2.8 (Corner / Non-corner) *A corner of a polygonal region, R , is a vertex of R with no internal angle larger than π . All other vertices of R are called non-corners of R .*

Remarks Suppose that the polygonal region, R , has a hole that itself can be treated as a simple polygon, H . Then the non-corners of H are corners of R .

Definition 2.9 (Side Chain) *A chain of edges of a (simple) polygon, P , between two consecutive corners of P is called a side chain of P .*

Note that the vertices of a side chain (without its end points) are exclusively non-corners. Of course, a side chain might consist of a single edge.

Definition 2.10 (Geodesic) *The shortest curve in a polygonal region, R , that connects two vertices, a and b , of R , is called the geodesic between a and b .*

Definition 2.11 (Pseudo-Triangle) *A simple polygon with exactly three corners is called a pseudo-triangle.*

Note that a triangle is also a pseudo-triangle. As already mentioned before, the convex hull, $\text{conv}(\nabla)$, of a pseudo-triangle, ∇ , is a triangle whose vertices are exactly the corners of ∇ . Furthermore, the geodesic between two corners of ∇ defines a side chain of ∇ and therefore lies entirely on ∇ 's boundary [3].

Definition 2.12 (Pseudo- k -Gon) *A simple polygon with exactly k corners is called a pseudo- k -gon. In particular, a simple polygon with 4 corners and 4 side chains is also called a pseudo-quadrilateral.*

Definition 2.13 (Pseudo-Triangulation) *Let R be a polygonal region. A pseudo-triangulation, $\mathcal{PT}(R)$, in R , is a cell complex in R whose cells are exclusively pseudo-triangles.*

Like for triangulations, a pseudo-triangulation of a finite point set S is the pseudo-triangulation, $\mathcal{PT}(R)$, in the polygonal region, R , where R equals $\text{conv}(S)$ and $\mathcal{PT}(R)$ uses exactly the points of S as vertices. Because triangles are also pseudo-triangles, triangulations are also pseudo-triangulations.

It is easy to see that faces of (full) triangulations have only one edge in common. However, the faces of pseudo-triangulations may intersect at two edges. It is well known that they cannot intersect in more than two edges.

Definition 2.14 (Double Adjacency) *Two pseudo-triangles, ∇_1, ∇_2 , of a pseudo-triangulation are in double adjacency if they have two edges in common.*

Observation 1 *In the case of a double adjacency, the union of the two pseudo-triangles is again a pseudo-triangle, [3].*

Definition 2.15 (Pointed / Non-pointed) *A vertex, v , of a pseudo-triangulation is called pointed if all its incident edges lie within an angle smaller than π . Otherwise, v is called non-pointed.*

In other words: v is pointed if and only if exactly one incident angle (angle between two consecutive incident edges of v) is larger than π .

Definition 2.16 (Pointed Pseudo-Triangulation) *A pseudo-triangulation, where each vertex is pointed, is called a pointed pseudo-triangulation.*

Note that all corners of a polygonal region, R , are always pointed. The more pointed vertices there are in $\mathcal{PT}(R)$, the less edges and faces it has. Furthermore, the *edge rank*, the difference between $|\text{vert}(\mathcal{PT})|$ and the number of pointed vertices of $\mathcal{PT}(R)$, is a minimum (zero) if $\mathcal{PT}(R)$ is a pointed pseudo-triangulation. As in this case also the number of edges of $\mathcal{PT}(R)$ is a minimum, pointed pseudo-triangulations are sometimes also referred to as *minimum pseudo-triangulations*. Be careful to not confuse minimum pseudo-triangulations with minimum weight (pseudo)-triangulations.

Definition 2.17 (Complete / Incomplete Vertex) *Consider a subset, B , of pseudo-triangles of a pseudo-triangulation. Let v be a vertex of a member of B . The vertex, v , is called complete in B , if v is corner in each of its incident pseudo-triangles of B . Otherwise, v is called incomplete in B .*

Note that B may be the entire pseudo-triangulation as well as a subset of disconnected pseudo-triangles. If we talk about completeness in entire pseudo-triangulations, we just say that a vertex is complete or incomplete, without specifying

the "subset". Further, note that the concepts of pointedness and completeness are rather similar. In the majority of cases pointed vertices are incomplete and complete vertices are non-pointed. The only difference occurs at vertices that are convex on the boundary of the pseudo-triangulation. Though such vertices are pointed they are also complete.

Observation 2 *Let v be a vertex of a pseudo-triangulation in a polygonal region R . If v is non-pointed then v is also complete. If v is pointed and not a corner of R then v is incomplete. If v is a corner of R then v is pointed and complete.*

Observation 3 *Consider a subset, B , of pseudo-triangles of a pseudo-triangulation, \mathcal{PT} . Let v be a vertex of a member of B . If v is complete in \mathcal{PT} , v is also complete in B . If v is incomplete in B , v is also incomplete in \mathcal{PT} .*

Definition 2.18 (Induced Polygonal Region of \mathcal{PT}) *Consider a pseudo-triangulation \mathcal{PT} with edge set E . Each polygonal region that can be formed by exclusively using edges of E is called an induced polygonal region of \mathcal{PT} .*

If we are talking about induced pseudo-triangles in this thesis, we mostly mean induced pseudo-triangles of \mathcal{PT} with at least one edge of \mathcal{PT} interior to them. Of course, also the pseudo-triangles that are forming \mathcal{PT} are induced pseudo-triangles of \mathcal{PT} , but we will refer to them just as pseudo-triangles of \mathcal{PT} .

2.3 Planar flipping

A so-called *flip* is an operation of constant combinatorial complexity, used to transform (pseudo-)triangulations into other ones [2, 10, 14, 18, 22].




Figure 2: Edge-exchanging flip types and non-flipable edge, [3].

The standard edge flip on triangulations is the so-called *Lawson flip* [17]. It takes two triangles, \triangle_1 and \triangle_2 , whose union is a convex quadrilateral and exchanges its diagonals. To generalize flipping to pseudo-triangulations, a different flip-definition is necessary, because the direct connections of the opposite corners of the union of two pseudo-triangles do not always exist in the union. The following definition (see for example [3]) uses a geodesic interpretation to make the Lawson flip suitable for pseudo-triangulations; see Figure 2(a) and (b) for examples.

Definition 2.19 (Edge-Exchanging Flip) *Consider two pseudo-triangles, ∇_1 and ∇_2 , of a pseudo-triangulation that are adjacent at edge e . Let g be the geodesic in $\nabla_1 \cup \nabla_2$ that connects the two corners, c_1 and c_2 , of ∇_1 and ∇_2 that are opposite to e . Moreover, g has to emanate from c_1 and c_2 within the angle spanned by the two incident edges of ∇_1 respectively ∇_2 . An edge-exchanging flip replaces e by the part, e' , of g interior to $\nabla_1 \cup \nabla_2$. If e and e' have one end-point in common, then the edge-exchanging flip is of non-crossing type, otherwise it is of crossing type.*

Remarks As the result of an edge-exchanging flip is (a part of) the geodesic between the two opposite corners of the flipped edge, an edge-exchanging flip is reversible. This means, if edge e' results from applying an edge-exchanging flip to edge e , immediately flipping edge e' results in edge e again. Of course, flipping edge e' can produce other results, if there were other flips applied to the pseudo-triangulation in the meantime.

Lemma 1 *An edge-exchanging flip does not alter the pointedness of any vertex.*

Proof. The edge e' (as well as e) is part of a geodesic between two corners of the pseudo-quadrilateral $\nabla_1 \cup \nabla_2$. As a geodesic uses only non-corners of $\nabla_1 \cup \nabla_2$ as additional vertices and connects them tangential, the pointedness of each vertex of $\nabla_1 \cup \nabla_2$ remains unchanged. \square

Lemma 2 *The common vertex v of a non-crossing edge-exchanging flip from e to e' has to be pointed. The simultaneous existence of both edges, e and e' , makes v non-pointed.*

Proof. The vertex v has to be a non-corner of the pseudo-quadrilateral, because it is part of both geodesics and thus is pointed. Now suppose that e and e' exist simultaneously and v is still pointed. The two edges of (the boundary of) the pseudo-quadrilateral incident to v are part of a side chain between 2 corners, c_1 and c_2 . The corner c_1 is connected to a different corner by a geodesic that has e as a part. The same applies for c_2 and e' . As v has to have its angle larger than π interior to the pseudo-quadrilateral, we will need a 5th corner. That is a contradiction, as a pseudo-quadrilateral has 4 corners, by Definition 2.12. \square

An edge-exchanging (or short exchanging) flip replaces one edge with another unique one and results in two valid pseudo-triangles. It is well known (see for example [3]) that in a pointed pseudo-triangulation each internal edge is flipable by an exchanging flip.

But observe that in the case of double-adjacent pseudo-triangles, ∇_1 and ∇_2 , there exists an ambiguous geodesic interpretation for an edge-exchanging flip, shown in Figure 3. Figure 3(a) shows the correct edge-exchanging flip while Figure 3(b) exemplifies a misinterpretation of the geodesics rule, where the geodesic nevertheless lies inside $\nabla_1 \cup \nabla_2$.




Figure 3: Ambiguous geodesics interpretation, [3].

Concerning triangulations, it is easy to see that if $\Delta_1 \cup \Delta_2$ is a non-convex quadrilateral, the edge e that is separating Δ_1 and Δ_2 is not flipable by an edge-exchanging (Lawson) flip, see Figure 2(c). In this case the non-corner of the non-convex quadrilateral is non-pointed. Interestingly, the non-convex quadrilateral, without edge e , forms a pseudo-triangle in this case.

If we exclusively operate on triangulations, we are not allowed to simply remove e . The cell that would remain after this operation, would not be a triangle. Therefore we would leave the class of triangulations. In contrast, removing e is a valid, combinatorial constant-size operation on pseudo-triangulations. This operation also has a geodesic interpretation that was used in [3] to introduce a novel type of edge flip, namely the *edge-removing flip*.

Definition 2.20 (Edge-Removing Flip, [3]) Consider ∇_1, ∇_2 and e as in Definition 2.19. The operation that removes e (without substitution) and results in a pseudo-triangle $\nabla_1 \cup \nabla_2$ is called *edge-removing flip*.

Lemma 3 An edge-removing flip alters the pointedness of exactly one vertex of the pseudo-triangulation from non-pointed to pointed, [3].

Lemma 4 Let \mathcal{PT} be a pseudo-triangulation of a polygonal region R . If an edge flip does not alter the pointedness of any vertex of \mathcal{PT} , also the completeness of all vertices remains unchanged. An edge flip that alters the pointedness of a vertex of \mathcal{PT} from non-pointed to pointed alters the completeness of this vertex from complete to incomplete.

Proof. As edges of R are incident to only one pseudo-triangle, edge flips are not applicable to them. By Observation 2 pointedness and completeness are analogous

concepts for vertices of \mathcal{PT} that are not corners of R . Further, vertices of \mathcal{PT} that are corners of R are always pointed and complete. \square

With the edge-exchanging and the edge-removing flip, each internal edge of an arbitrary pseudo-triangulation is flipable. Further the edge-removing flip, and its inverse, the edge-inserting flip, enables us to shift between pseudo-triangulations of different edge rank. Moreover, based on the new flip type, in [3] another operation was specified that enables us to delete certain vertices from a pseudo-triangulation in a well-defined way that guarantees a pseudo-triangulation as result.




Figure 4: Vertex-removing flip, [3].

Consider a vertex, v , of a pseudo-triangulation, \mathcal{PT} , with two incident edges, e_1 and e_2 , see Figure 4 (e_1, e_2 are bold). The vertex v and its incident edges partition an induced pseudo-triangle, ∇ , of \mathcal{PT} into two other pseudo-triangles of \mathcal{PT} that are in double-adjacency (at e_1 and e_2). We can remove v by applying an edge-removing flip to both its incident edges simultaneously. This leaves ∇ as empty pseudo-triangle. Therefore a valid pseudo-triangulation remains. Alternatively, we could apply the ambiguous geodesics interpretation (Figure 3) to e_1 and e_2 in succession. Thereby we would produce the same result.

Definition 2.21 (Vertex-Removing Flip, [3]) *The operation which deletes a vertex of degree 2 of a pseudo-triangulation along with its two incident edges and thereby joins two double-adjacent pseudo-triangles to one, is called a vertex-removing flip.*

Definition 2.22 (Flipping Pair) *Let e' be the edge that results from edge e in an edge-exchanging flip. Then we call the edges e and e' a flipping pair.*

Definition 2.23 (Potential Flipping Pair) *Assume that the edges e and e' exist at the same time in a pseudo-triangulation and that it is possible to delete either edge with an edge-removing flip. Then we call e and e' a potential flipping pair if they would be a flipping pair if either e or e' would not exist.*

Remarks It is only possible to remove one of this edges. After deleting one edge it is not possible to delete the other. Otherwise they would not form a flipping pair if one is absent. Moreover, in that case, they would form a flipping pair of a non-crossing edge-exchanging flip and thus would have one vertex in common.

Observation 4 *The common vertex, v , of a potential flipping pair, e, e' , (Definition 2.23) is non-pointed before, and pointed after removing one edge, e or e' , with an edge-removing flip.*

Proof. As e and e' exist at the same time, they can only be a flipping pair for a non-crossing edge-exchanging flip if one of them has been removed. By Lemma 2, the common vertex of a non-crossing edge-exchanging flip is pointed, and the simultaneous existence of e and e' makes v non-pointed. \square

Definition 2.24 (Inner Tangents) *Let R be a polygonal region. A line segment, l , between two vertices of R and in the interior of R is called an inner tangent for R if l is part of some pointed pseudo-triangulation in R .*

Lemma 5 *Let C be a convex polygon interior to a pseudo-triangle ∇ . There exist exactly 6 inner tangents for the polygonal region $R = \nabla \setminus C$. There exist combinations of exactly 3 inner tangents that form a pointed pseudo-triangulation in R .*

Proof. All non-corners of R (non-corners of ∇ and corners of C) are pointed. To maintain pointedness, the only allowed edges are geodesics inside R from each corner of ∇ that are tangents to C . As each corner sees a left and a right side of C there exist exactly 2 inner tangents from a corner of ∇ to a vertex of C . Since ∇ has 3 corners and C is not the empty set, this gives 6 possible edges.

The number of necessary and sufficient edges to obtain a pointed pseudo-triangulation is reduced from 6 to 3 because at each time 2 of the 6 edges are flipping pairs (Definition 2.22) and therefore are mutually exclusive: A flipping pair is either crossing or has one vertex in common. In the former case it is obvious that the two edges are mutual exclusive. In the latter case the existence of both edges violates the pointedness of their common vertex, by Observation 4.

Three inner tangents partition R into three cells and the boundary of ∇ into three "pieces", the side chains of ∇ . Furthermore, three inner tangents split the boundary of C into three concave chains (viewed from inside R). Along with the three inner tangents themselves, this makes 9 side chains parted by 9 corners altogether. As we got three cells, each of them can have three side chains and three corners. As every vertex of R is pointed, this describes a pointed pseudo-triangulation with three pseudo-triangles. \square

The following lemma has been established and proved by M.Pocchiola and G.Vegter in [18]. The formalism has been slightly adapted, since their paper deals with disjoint convex obstacles. The pseudo-triangulations they construct are pointed, although not explicitly termed as such.

Lemma 6 (Lemma 3 from [18]) *The number of tangents in a pointed pseudo-triangulation of a collection of n disjoint convex objects is $3n - 3$.*

Using Lemmata 5 and 6 we can provide the exact number of inner tangents necessary for a pointed pseudo-triangulation within a pseudo-triangle with an arbitrary number of disjoint convex holes.

Lemma 7 *The number of inner tangents in a pointed pseudo-triangulation of a collection of n disjoint convex polygons surrounded by a pseudo-triangle is $3n$.*

Proof. From Lemma 6 we know that we need $3n - 3$ tangents (or edges) for a pointed pseudo-triangulation of n disjoint convex objects. This pseudo-triangulation contains the edges of the convex hull, CH , of all objects. Thus, we gain a single convex object CH inside a pseudo-triangle ∇ . By Lemma 5 it takes another 3 tangents to pseudo-triangulate the polygonal region $\nabla \setminus CH$ in a pointed way. Therefore $3n - 3 + 3 = 3n$ tangents are needed.

Note that there exist a lot of different pointed pseudo-triangulations for ∇ and its internal disjoint convex polygons. (Even some, where no single edge of CH is included.) As two pointed pseudo-triangulations can be transformed into each other by applying $O(n \log^2 n)$ edge-exchanging flips [3], every possible pointed pseudo-triangulation can be created by flipping edges of the above obtained pointed pseudo-triangulation. As edge-exchanging flips do not change the number of edges, the number of inner tangents stays the same. \square

2.4 Surfaces for pseudo-triangulations

In [3] it is proved that pseudo-triangulations have realizations as *polyhedral surfaces* in three-space and showed how to construct them. Further an attribute is introduced for pseudo-triangulations called *projectivity*, a *surface interpretation of flips* is presented, and it is proved that *flipping to optimality* is always possible within the class of pseudo-triangulations. As this thesis is mainly based on [3], this section is more or less a synopsis from that work. Nevertheless, there are slight differences to adapt to the formalism and layout of this thesis.

Definition 2.25 (Polyhedral Surface) *A polyhedral surface is the graph of a continuous and piecewise-linear function whose domain is a polygonal region.*

Definition 2.26 (Convex, Reflex and Planar Edge) *Let e be an edge in a polyhedral surface. The edge e is called convex if there exists a line segment that intersects e at exactly one interior point and everywhere else lies below the surface. The edge e is called reflex if there exists a line segment that intersects e at exactly one interior point and everywhere else lies above the surface. If e is neither convex nor reflex then we call e planar.*

Theorem 1 (Surface Theorem, [3]) *Let R be a polygonal region, and let \mathcal{PT} be any pseudo-triangulation within R . Let \mathbf{h} be a vector assigning a height to each complete vertex of \mathcal{PT} . For each choice of \mathbf{h} , there exists a unique polyhedral surface $\mathcal{F}(\mathcal{PT}, \mathbf{h})$ above the domain R , that respects \mathbf{h} and whose edges project vertically to (a subset of) the edges of \mathcal{PT} .*

This theorem holds for arbitrary polygonal regions, including such with holes. This will turn to be a fundamental property for Section 6. Further the theorem admits surfaces whose edges project only to a subset of the edges of the underlying pseudo-triangulation. This happens because planar edges do not project vertically.

Though there exist various different surfaces for a given pseudo-triangulation, depending on the choice of the height vector \mathbf{h} , there exist "special" pseudo-triangulations that have no surface realization whose edges project back to exactly the original edges. To classify such pseudo-triangulations, [3] introduces the concept of *projectivity* for pseudo-triangulations, respectively cell complexes in general.

Definition 2.27 (Projectivity) *A cell complex \mathcal{C} is called projective if there exists some polyhedral surface whose set of edges projects exactly to the set of edges of \mathcal{C} . Otherwise, we call \mathcal{C} non-projective.*

This property highly depends on the geometric embedding of the pseudo-triangulation. An "epsilon small" change of the planar vertex set can change the property. This is an undesired effect, as we want to stay as independent as possible from the geometric realization. Therefore we introduce an advanced version of projectivity, called *combinatorial projectivity*. There will be some properties later that also will be combinatorially defined for the same reason, see the Definitions 2.32, 2.33, 2.34. Combinatorial in this context means that the property does not change if the underlying vertex set is perturbed by an arbitrarily small $\epsilon > 0$.

Definition 2.28 (S_ϵ) *The vertex set S_ϵ is an ϵ -perturbation of S if the vertices of S_ϵ are the vertices of S perturbed by some arbitrary small $\epsilon > 0$ such that the order types ($[4]$, $[5]$, $[11]$, $[16]$) of S_ϵ and S are equal.*

According to the ϵ -perturbed vertex set S_ϵ we want to define an ϵ -perturbed pseudo-triangulation that rules out any special geometric realizations, like the one shown in Figure 5.

Definition 2.29 (\mathcal{PT}_ϵ) *Let \mathcal{PT} be a pseudo-triangulation with vertex set S . Further, let S_ϵ be an ϵ -perturbation of S . The pseudo-triangulation \mathcal{PT}_ϵ that is combinatorial equivalent to \mathcal{PT} and has S_ϵ as its vertex set is called an ϵ -perturbation of \mathcal{PT} .*

Definition 2.30 (Combinatorial Projectivity) *Consider R and \mathcal{PT} as in Theorem 1. Further, let S be the set of vertices of \mathcal{PT} . Then \mathcal{PT} is called combinatorial projective, if there exists some ϵ -perturbation S_ϵ of S , such that \mathcal{PT}_ϵ is projective. Otherwise, we call \mathcal{PT} combinatorial non-projective.*

Observe the trivial fact that projectivity implies combinatorial projectivity and combinatorial non-projectivity implies non-projectivity.

Figure 5 exemplifies the difference between projectivity and combinatorial projectivity. The bold edge of the pseudo-triangulation \mathcal{PT} in Figure 5(a) is planar

in each surface $\mathcal{F}(\mathcal{PT}, \mathbf{h})$, regardless of the height vector \mathbf{h} . This is because the intersection of the prolongations (dotted edges) lies exactly on another edge. Therefore only the pseudo-triangulation in Figure 5(b) is projective, because one vertex (the lower right one) was moved by an arbitrary small ϵ . However, both pseudo-triangulations are combinatorial projective. For combinatorial non-projective examples see Figures 10(a+b) on page 21 or Figures 14(a+b) on page 34.




Figure 5: The difference between the concepts of projectivity and combinatorial projectivity.

Though we have banned geometric influence from the 2-dimensional vertex set by using \mathcal{PT}_ϵ , it is possible that the surface $\mathcal{F}(\mathcal{PT}_\epsilon, \mathbf{h})$ fails to be strictly convex or reflex for some edges of \mathcal{PT}_ϵ . This happens if either \mathcal{PT}_ϵ is (combinatorial) non-projective or the height vector, \mathbf{h} , is degenerate. To foreclose the latter case, [3] introduces another term.

Definition 2.31 (Generic Height Vector) *A height vector, \mathbf{h} , is called generic, if \mathbf{h} witnesses the (combinatorial) projectivity of \mathcal{PT}_ϵ .*

Now the only reason for a surface, $\mathcal{F}(\mathcal{PT}_\epsilon)$, whose set of edges is not projecting exactly to the edges of \mathcal{PT} , is that \mathcal{PT} is combinatorial non-projective, for example (but not limited to) the case of a double-adjacency in \mathcal{PT} . Before stating a definition from [3] that was designed to decide (combinatorial) projectivity by only examining the (2-dimensional) pseudo-triangulation, we want to introduce some new definitions and observations that discuss the connectivity between 3-space properties of or within the surface and the combinatorial projectivity of the underlying pseudo-triangulation.

Definition 2.32 (Combinatorial Coplanarity) *A subset of pseudo-triangles of a pseudo-triangulation \mathcal{PT} is called combinatorial coplanar if their corresponding facets all lie in a common plane, for all height vectors \mathbf{h} , and all ϵ -perturbations of \mathcal{PT} .*

Remarks Note that Definition 2.30 and Definition 2.32 are not equivalent. Figures 10(a+b) on page 21 show pseudo-triangulations that are combinatorial non-projective and contain combinatorial coplanarities (gray pseudo-triangles). But Figure 10(c) contains combinatorial coplanarities, in spite of being combinatorial projective.

Observation 5 *Pseudo-triangles that have 3 vertices in common are combinatorial coplanar.*

Note that here it is important that we have no collinear points! Observe further that the converse of this statement is not true. In the examples in Figures 14(a+b) on page 34 no pseudo-triangle has 3 vertices in common with any other pseudo-triangle. Anyhow, the gray set of pseudo-triangles in Figure 14(b) is combinatorial coplanar. In Figure 14(a) even the whole pseudo-triangulation is combinatorial coplanar

Definition 2.33 (Combinatorial Planar Edge) *An edge e of the pseudo-triangulation \mathcal{PT} is called combinatorial planar if e is incident to 2 pseudo-triangles that are combinatorial coplanar. Edge e is called combinatorial non-planar, otherwise.*

Definition 2.34 (Combinatorial Planar Vertex) *A vertex of a pseudo-triangulation \mathcal{PT} is called combinatorial planar if all its incident edges in \mathcal{PT} are combinatorial planar.*

It is trivial to see that all edges incident to a combinatorial planar vertex are combinatorial planar edges. But observe that there may exist combinatorial planar edges that are not incident to any combinatorial planar vertex.

Observation 6 *The set of all pseudo-triangles incident to a vertex v of a pseudo-triangulation is combinatorial coplanar iff v is combinatorial planar.*

Observation 7 *The existence of a combinatorial planar vertex or a combinatorial planar edge implies that the pseudo-triangulation is combinatorial non-projective.*

As the (combinatorial) projectivity is a property of the 2-dimensional pseudo-triangulation, we do not want to generate a surface to decide it, not only because of the possibility of degenerate height vector. Therefore [3] introduces the concept of *stability* that was meant to decide projectivity by only examining (subsets of) the pseudo-triangulation.

Definition 2.35 (Stability, [3]) *Let \mathcal{PT} be a pseudo-triangulation in a simple polygon P and let S be the vertex set of \mathcal{PT} (with $\text{vert}(P) \subset S$). Then \mathcal{PT} is called stable if no subset of incomplete vertices of S can be eliminated, along with their incident edges such that:*

- (1) *a valid pseudo-triangulation \mathcal{PT}' remains, and*
- (2) *the status of each vertex of \mathcal{PT}' is the same as in \mathcal{PT} .*

In other words: \mathcal{PT} is stable if no induced pseudo-triangle, ∇ , of \mathcal{PT} exists such that every non-corner of ∇ and each vertex inside ∇ is incomplete. Observe that double-adjacencies are ruled out in a stable pseudo-triangulation. Moreover,

also pointed pseudo-triangulations on polygons P , where $\partial(P)$ is a pseudo-triangle, are not stable.

To decide combinatorial projectivity by deciding stability a theorem was stated in [3] that connected both concepts in the way that a pseudo-triangulation is combinatorial projective if and only if it is stable. This turned out to be wrong, as shown in the counter example in Figure 14(b) on page 34 that is stable but combinatorial non-projective. Figure 10(b) on page 21 shows another counter example, where the pseudo-triangulation is also stable but combinatorial non-projective. Nevertheless, for at least some examples concluding from non-stable to combinatorial non-projective is correct, see Figures 14(a) and 10(a).

This leads to the core task of this thesis, namely to find out when exactly stable pseudo-triangulations are combinatorial non-projective, to classify those cases and to discover a way to fix the stability concept. This turned out to be more complex than thought and thus the rest of this thesis, after this section, deals with this problem. Therefore, for now, we do not go into detail, as the remaining surface introduction does not use this concept and therefore stays correct.

To be able to transform surfaces into others, [3] provides surface interpretations of the planar flips introduced in Section 2.3.

Definition 2.36 (Surface Flip) *Let \mathcal{PT} be some combinatorial projective pseudo-triangulation in a polygonal region R and let \mathcal{PT}' be the unique pseudo-triangulation obtained from \mathcal{PT} by applying a single admissible flip. Further assign height vectors \mathbf{h} and \mathbf{h}' to \mathcal{PT} and \mathcal{PT}' such that \mathbf{h} is generic and coincides with \mathbf{h}' for all vertices being complete in both structures. By Theorem 1, there exist 2 unique surfaces $\mathcal{F}(\mathcal{PT}, \mathbf{h})$ and $\mathcal{F}'(\mathcal{PT}', \mathbf{h}')$. An operation that transforms $\mathcal{F}(\mathcal{PT}, \mathbf{h})$ into $\mathcal{F}'(\mathcal{PT}', \mathbf{h}')$ is called a surface flip.*

Surface flips cause local and constant-size changes in a combinatorial sense. But geometrically they can cause global changes, as a single flip can change vertex heights of many vertices.

Our main task, when flipping in surfaces, is to reach the *maximal locally convex surface* which, by Definition 2.42, does not contain any reflex edges. Therefore we are only interested in surface flips applied to reflex edges. Thus we define:

Definition 2.37 (Improving Surface Flip) *A surface flip applied to a reflex edge in a surface is called an improving surface flip.*

Both, the edge-exchanging (Definition 2.19) and the edge-removing flip (Definition 2.20) have an interpretation as improving surface flips. According to Definition 2.36, [3] defines a *convexifying flip* and the *planarizing flip*.

Definition 2.38 (Convexifying Flip) *An improving surface flip that corresponds to an edge-exchanging flip in the plane, is called convexifying flip.*




Figure 6: Two convexifying flips, [3].




Figure 7: A planarizing flip, [3].

A convexifying flip exchanges a reflex edge with a convex one, hence the name. See Figure 6 for examples. The properties for the edge-exchanging flip, Lemmata 1 and 2, also hold for convexifying flips.

Definition 2.39 (Planarizing Flip) *An improving surface flip that corresponds to an edge-removing flip in the plane, is called planarizing flip.*

The planarizing flip removes one reflex edge from the surface. For an example of the geometric change, see Figure 7. Thereby two facets, previously adjacent at the flipped reflex edge, are joined to one single facet that flattens out. Also the planarizing flip keeps the property, Lemma 3, from its analogon.

A convexifying or planarizing flip may result in a pseudo-triangulation that is not combinatorial projective any more. As in [3] combinatorial projective is identical to stable, what turned out to be incorrect, the *trivial flip* is designed in [3] to eliminate the edges and vertices from the surface that stem from the subset of incomplete vertices that causes the non-stability, see Definition 2.35.

Nevertheless, the trivial flip turned out to be very useful, see Section 6. Thus we introduce a slightly adapted version of the trivial flip definition:

Definition 2.40 (Trivial Flip) Let \mathcal{PT} be some pseudo-triangulation. Let S_p be the subset of all combinatorial planar vertices of \mathcal{PT} . Eliminating each vertex $v \in S_p$ by applying a sequence of vertex-degree-reducing edge-exchanging and edge-removing flips, followed by a degree-2 vertex-removing flip to v , is called a trivial flip.

A trivial flip can be applied after each (improving) surface flip to delete all combinatorial planar vertices within the new surface. All edge flips in the sequence of a trivial flip are applied to combinatorial planar edges and either delete one edge (edge-removing) or exchange one edge with another combinatorial planar edge (edge-exchanging).

Lemma 8 During a trivial flip, edge-exchanging flips create only combinatorial planar edges.

Proof. All facets incident to a combinatorial planar vertex v lie in a common plane by Observation 6. Inside this set, B , of combinatorial coplanar pseudo-triangles all edges are combinatorial planar by Definition 2.33. A degree reducing edge-exchanging flip replaces an edge incident to v with one that stays within B . Thus this edge is also combinatorial planar. \square




Figure 8: Example where each combinatorial planar edge is incident to some combinatorial planar vertex.

Remarks Even if each combinatorial planar edge in the surface is incident to at least one combinatorial planar vertex, combinatorial planar edges may remain after a trivial flip. See Figure 8 for an example.

As already mentioned earlier, one task in flipping in surfaces is to reach convexity, which equals the lower convex hull, if the underlying polygonal domain is a convex polygon. In [3] an optimization problem was stated, which was solved with the optimality theorem and a second theorem which are restated as one theorem below:

Definition 2.41 (Locally Convex Function) Let P be a polygon in the plane. A real-valued function f with domain P is called locally convex if f is convex on each line segment internal to P .

Optimization Problem:

Let S be a set of vertices within some simple polygon, P , with $\text{vert}(P) \subset S$. Let \mathbf{h} be a vector assigning some real value to each vertex in S . Given P , S and \mathbf{h} , find the maximal locally convex function f^* on the domain P which fulfills $f^*(v_i) \leq h_i$ for each $v_i \in S$.

Theorem 2 (Optimality Theorem, [3]) Let S be a set of vertices within some simple polygon, P , with $\text{vert}(P) \subset S$ and let \mathbf{h} be a height vector for S . Let f^* be the unique maximal locally convex function on P that is bounded from above by \mathbf{h} .

- (1) The graph \mathcal{F}^* of f^* projects to a pseudo-triangulation in the generic case.
- (2) \mathcal{F}^* can be constructed from any triangular surface \mathcal{F} on P for S and \mathbf{h} , by applying any improving surface flip (and trivial flip) sequence of sufficient, but finite, length.

To simplify matters we state an additional definition for the graph of the maximal locally convex function:

Definition 2.42 (Maximal Locally Convex Surface) We call the graph \mathcal{F}^* of f^* from Theorem 2 the maximal locally convex surface, respectively optimal surface for short.




Figure 9: Local and global optima, [3].

Observe that it is not always possible to reach the maximal locally convex surface within the class of triangulations, see Figure 9. The numbers at the vertices denote

their heights and reflex edges are bold. Flipping one of the two reflex edges in the initial surface in Figure 9(a) leads to either the triangular surface in Figure 9(b) or Figure 9(c). As in both examples the remaining reflex edge is not flipable without leaving the class of triangulations, these are local optima. The global optimum \mathcal{F}^* , shown in Figure 9(d), is a pseudo-triangulation.

Finally we want to combine two characteristics of surfaces that stem from the optimality theorem and were mentioned in [3] in Section 7.1 "flipping to optimality", to a lemma.

Lemma 9 (Section 7.1 in [3]) *In the optimal surface there exists no internal pointed vertex. Each internal pointed vertex of a surface has at least one incident reflex edge.*

3 3-Reducibility

3.1 The problem

Up to now double adjacencies were characterized by two pseudo-triangles sharing two edges (Definition 2.14). These edges are combinatorial planar because the two facets reside in the same plane. But the reason for this circumstance is that pseudo-triangles sharing two edges automatically share three vertices¹ which define a plane where both facets must reside. Interestingly, it is possible to pseudo-triangulate a point set in such a way that two pseudo-triangles share three vertices without sharing 2 edges. Certainly also in this case the two facets share the same plane and therefore the separating edge (if one exists) is combinatorial planar.

See Figure 10 for an example: Bold edges are combinatorial planar and gray pseudo-triangles are combinatorial coplanar. Though all three examples show two combinatorial coplanar facets, they differ in the number of combinatorial planar edges. Figure 10(a) shows a classical double adjacency whereas Figures 10(b)+(c) show "deformed" examples.




Figure 10: (a) shows a "normal" double adjacency. (b) and (c) show deformed double adjacencies where in (b) exists one combinatorial planar edge, e , and in (c) none.

Considering the stable definition for pseudo-triangulations we observe, though such a *deformed double adjacency* (Figure 10(b)+(c)) is stable, it is combinatorial non-projective, if there exists a separating edge, e (Figure 10(b)). That violates the statement that a pseudo-triangulation \mathcal{PT} , with vertex set S , can be made projective by perturbing S by some arbitrarily small ϵ , if it is stable [3]. Furthermore, a trivial flip will not remove e without substitution. We also observe, that the projection of a surface \mathcal{F} of \mathcal{PT} does no longer belong to the class of pseudo-triangulations for any (generic) height vector. On this account we have to take a closer look to this kind of edges in pseudo-triangulations that are *hidden* in the projections of their surfaces.

¹Let us stress the fact again, that there are no collinear vertices!

3.2 3-Reducibility

First of all it is important to acquire a criterion that allows us to locate all combinatorial planar edges within a given pseudo-triangulation. Consider a pseudo-triangulation \mathcal{PT} . See Figure 11 for three examples. According to the *Surface Theorem* [3] the height of an incomplete vertex, v , of \mathcal{PT} depends on the heights of the corners of its unique corresponding pseudo-triangle of \mathcal{PT} where v is a non-corner. If one of these corners, c , is incomplete, its height again depends on the heights of three vertices. Thus the height of v is depending on the heights of at most five other vertices after the first step, and so on. Figure 11(a) illustrates an example where the height of v depends on 5 other vertices (square dots) whereas in the examples in Figure 11(b) and (c) the height of v depends on 4 resp. 3 other vertices.




Figure 11: Examples of pseudo-triangulations where the height of an incomplete vertex v depends on the height of 5 (figure (a)), 4 (figure (b)) or 3 (figure (c)) other vertices (square dots).

Definition 3.1 (Height Defining Vertex) Let B be a subset of pseudo-triangles of a pseudo-triangulation \mathcal{PT} . Let v be a vertex of some member of B . We call v height defining (for B) if v is complete in B .

Observation 8 Consider B and \mathcal{PT} as in Definition 3.1. A vertex of $\text{conv}(\bigcup B)$ is complete in B and a vertex that is incomplete in B has to reside within $\text{conv}(\bigcup B)$.

Proof. A vertex that is corner of the boundary of $\bigcup B$ is complete in B , by Observation 2. As $\text{conv}(\bigcup B)$ is a boundary of $\bigcup B$ on which each vertex is a corner, every vertex on $\text{conv}(\bigcup B)$ is complete in B . From this it follows that vertices of members of B that are incomplete in B , cannot be on $\text{conv}(\bigcup B)$ and thus have to reside within $\text{conv}(\bigcup B)$. \square

Note that, by applying Observation 3, vertices that are complete in \mathcal{PT} , are height defining vertices for each subset of pseudo-triangles of \mathcal{PT} , where they are a vertex of at least one pseudo-triangle. On the other hand, height defining vertices for some subset of pseudo-triangles of \mathcal{PT} are not necessarily complete in \mathcal{PT} .

Definition 3.2 ((Height) Back-Tracing) Consider the subset B as in Definition 3.1. Let v be height defining for B , but incomplete in \mathcal{PT} . Adding to B the unique pseudo-triangle ∇_v where v is non-corner is called a (height) back-tracing step for v .

In $B \cup \nabla_v$, the vertex v is not complete any more. Therefore v is no longer height defining. Moreover, the corners of ∇_v , c_1, c_2, c_3 , are not necessarily height defining for $B \cup \nabla_v$. In particular, corner c_i is not height defining if c_i already belonged to B and was not height defining there.

Remarks The interested reader may have noticed that we can alternatively look at the pointedness of a vertex that is complete in the actual subset of pseudo-triangles, to decide whether it can be back-traced or not. Nevertheless, the concept of completeness is sufficient and using another concept may cause confusion.

Definition 3.3 (3-Reducibility) Consider a subset B_0 of pseudo-triangles of \mathcal{PT} . Let $\mathcal{B}(B_0)$ be the collection of all subsets $B \subseteq \mathcal{PT}$ such that there exists a sequence of height back-tracing steps that transforms B_0 into B . That is

$$\mathcal{B}(B_0) = \{B \mid B \text{ can be generated from } B_0 \text{ by back-tracing}\}.$$

We call B_0 3-reducible if $\mathcal{B}(B_0)$ contains a member with exactly 3 height defining vertices.

Remarks Also applying no back-tracing step is a valid sequence of back-tracing steps. This means that B_0 itself is member of $\mathcal{B}(B_0)$. This implicates that a subset with only 3 height defining vertices is obviously 3-reducible.

By the *Surface Theorem*, the height of each incomplete vertex v of \mathcal{PT} is described by a linear equation containing the heights of v and the corners c_1, c_2, c_3 of its unique pseudo-triangle ∇_v . We symbolically abbreviate this relation as $v = f(c_1, c_2, c_3)$.

Using this notation we are able to rewrite the Definitions 3.1, 3.2 and 3.3 to linear systems of equations. Look back at the examples from Figure 11 to get a hint of how to do that.

In Figure 11 we start with a single vertex v . B is empty, so v is complete in B . Thus v is height defining for B , abbreviated as $v = f(v)$. The first back-tracing step now replaces $v = f(v)$ with $v = f(a, b, c)$ and additionally adds $a = f(a)$, $b = f(b)$ and $c = f(c)$. For the linear system this represents the replacement of 1 degree of freedom with 3 new ones, namely a, b, c . The second step results in a different number of degrees of freedom for the 3 different examples. For simplification let us name the corners of the corresponding pseudo-triangle ∇_c of c with c_1, c_2, c_3 . Then the equation $c = f(c_1, c_2, c_3)$ will replace $c = f(c)$ in the linear system in the next back-tracing step. For Figure 11(a) additionally the relations $c_1 = f(c_1)$, $c_2 = f(c_2)$ and $c_3 = f(c_3)$ will be added. This means that one degree of freedom, c , has been replaced by 3 new ones, c_1, c_2, c_3 , what makes 5 altogether. In Figure 11(b) one corner, lets say c_1 , is identical to v . As there already exists an equation for v in the

linear system, only $c_2 = f(c_2)$ and $c_3 = f(c_3)$ will be added. Thus there remains a linear system with 4 degrees of freedom. And in Figure 11(c) also a second corner, c_2 , equals a vertex, b , that is already part of the linear system which results in only 3 degrees of freedom in total, since we get only $c_3 = f(c_3)$ in addition.

Let us describe more formally the commonalities of the previously stated definitions and linear systems of equations.

Definition 3.4 ($\mathcal{LS}(B)$) *Let B be a subset of \mathcal{PT} as in Definition 3.1. $\mathcal{LS}(B)$ is called the corresponding linear system of B if it contains the height dependency equations:*

$$\begin{aligned} v = f(c_1, c_2, c_3) & \quad \dots \text{ for each } v \text{ being incomplete in } B \text{ with unique } \nabla_v \\ v = f(v) & \quad \dots \text{ for each } v \text{ being complete in } B \end{aligned}$$

Observation 9 *Consider B and $\mathcal{LS}(B)$ as in Definition 3.4. Then we can state:*

- (1) *The number of height defining vertices for B (see Definition 3.1) equals the degrees of freedom of $\mathcal{LS}(B)$.*
- (2) *Height back-tracing of a vertex v of B (see Definition 3.2) is equivalent to replacing $v = f(v)$ with $v = f(c_1, c_2, c_3)$ in $\mathcal{LS}(B)$ and adding $c_i = f(c_i)$ to $\mathcal{LS}(B)$ if c_i is complete in B .*
- (3) *If B_0 is 3-reducible (see Definition 3.3), $\mathcal{B}(B_0)$ contains a member B such that $\mathcal{LS}(B)$ has exactly 3 degrees of freedom.*

Using linear systems of equations for an algorithm to decide 3-reducibility would require to remember lots of sets of equations. To simplify matters we define a marked vertex set $I(B)$ as a representation for the linear system $\mathcal{LS}(B)$. The goal is to be able to decide 3-reducibility by only remembering these marked vertex sets, without storing any linear system or pseudo-triangulation subset during the algorithm. $I(B)$ will contain certain vertices for B and a flag for each vertex.

Definition 3.5 (Marked Vertex Set $I(B)$) *Consider B and $\mathcal{LS}(B)$ as in Definition 3.4. The so-called "marked vertex set for B ", $I(B)$, contains each vertex that is corner for at least one pseudo-triangle in B . In addition, each vertex in $I(B)$ is "marked" if it is not height defining for B and remains "unmarked" otherwise.*

By this definition, vertices that are a corner in no pseudo-triangle of B do not belong to $I(B)$. The following lemma (and its proof) shows that $I(B)$ nevertheless is a data structure providing sufficient information to decide 3-reducibility of a subset B .

Lemma 10 *Consider B and $\mathcal{LS}(B)$ as in Definition 3.4, and $I(B)$ as in Definition 3.5. Then the number of unmarked vertices in $I(B)$ equals the number of degrees of freedom of $\mathcal{LS}(B)$.*

Proof. By Observation 9, the number of degrees of freedom of $\mathcal{LS}(B)$ equals the number of height defining vertices for B . It remains to be shown that vertices of pseudo-triangles of B that are not in $I(B)$ cannot be height defining for B . Such a vertex cannot be corner for any pseudo-triangle of B , by Definition 3.5, and thus belongs to a unique pseudo-triangle where it is non-corner. Therefore it is not height defining for B (or any subset of pseudo-triangles of B). \square


It remains to define how to execute a back-tracing step using $I(B)$. A vertex $v \in I(B)$ can be back-traced if v is *unmarked* in $I(B)$ and incomplete in \mathcal{PT} . The back-tracing step is executed by *marking* v and adding the corners of ∇_v to $I(B)$ (*unmarked*) if not already in $I(B)$. Observe that in the back-tracing step only corners are added to $I(B)$.

Before stating a decision algorithm for 3-reducibility, we give two short examples in Figure 12. Incomplete vertices are labeled with numbers, complete ones with capitals. Vertices that are *marked* in $I(B)$ are underlined in the vertex set representation. The two examples run on identical pseudo-triangulations but with different initial subsets B_0 , showed as gray pseudo-triangles. For better understanding of the relation between $\mathcal{LS}(B)$ and $I(B)$, the state of $\mathcal{LS}(B)$ is displayed for each step, as well.


Both examples start with an initial subset B_0 containing 2 pseudo-triangles. In the first one (Figure 12(a)) $I(B_0)$ contains only one height defining vertex that is incomplete in \mathcal{PT} . Therefore only one back-tracing step is possible — adding the equation $2 = f(A, B, 3)$. This results in a vertex set $I(B_1)$ containing 4 vertices that are complete in \mathcal{PT} . The corresponding linear system has 5 degrees of freedom, indicated by the 5 vertices that are not underlined in $I(B_1)$. Again we have a single possibility for back-tracing — adding the equation $3 = f(B, C, 1)$. We get a vertex set, $I(B_2)$, containing the same 4 complete vertices as $I(B_1)$. But this time there also only exist 4 degrees of freedom and thus no further back-tracing is possible. Therefore B_0 is not 3-reducible since we created all possible $B \in \mathcal{B}(B_0)$.

Also the second example starts with a single possibility of back-tracing. Observe, that vertex 1 is not part of $I(B_0)$, because it is no corner for B_0 . Adding the equation $3 = f(B, C, 1)$ results in a marked vertex set, $I(B_1)$, containing 3 vertices complete in \mathcal{PT} and 4 vertices that are height defining for B_1 . Back-tracing vertex 1, does not add any additional vertex to the marked vertex set, but it makes vertex 1 *marked* in $I(B_2)$. Only 3 vertices remain *unmarked* in $I(B_2)$. These are the 3 vertices of B_2 that are complete in \mathcal{PT} and also height defining for B_2 . Thus the initial B_0 is 3-reducible.

The interested reader may have noticed that the approach of adding symbolical height equations to $\mathcal{LS}(B)$ differs from adding a pseudo-triangle ∇ to B , in the sense that adding ∇ may also add additional non-corners not belonging to $I(B \cup \nabla)$. Nevertheless, this does not change the degrees of freedom by Lemma 10. If such vertices occur as degrees of freedom in later back-tracing steps, they can be back-traced further with one additional step to three vertices that are already members of the visited vertex set. For example, this happened in the last back-tracing step



(a) $\dots B_0 = \{\{A, 2, 1, C\}, \{1, 2, D\}\}$



(b) $\dots B_0 = \{\{A, 2, 1, C\}, \{A, B, 3, 2\}\}$

<p>initial subset:</p> $I(B_0) : \{\underline{1}, 2, A, C, D\}$ $\mathcal{LS}(B_0) : \left\{ \begin{array}{l} A = f(A) \\ C = f(C) \\ D = f(D) \end{array} \right \begin{array}{l} 2 = f(2) \\ 1 = f(A, 2, C) \end{array}$ <p>\Rightarrow 4 degrees of freedom</p> <p>back-tracing: $2 = f(A, B, 3)$</p> $I(B_1) : \{\underline{1}, \underline{2}, 3, A, B, C, D\}$ $\mathcal{LS}(B_1) : \left\{ \begin{array}{l} A = f(A) \\ B = f(B) \\ C = f(C) \\ D = f(D) \end{array} \right \begin{array}{l} 3 = f(3) \\ 1 = f(A, 2, C) \\ 2 = f(A, B, 3) \end{array}$ <p>\Rightarrow 5 degrees of freedom</p> <p>back-tracing: $3 = f(B, C, 1)$</p> $I(B_2) : \{\underline{1}, \underline{2}, \underline{3}, A, B, C, D\}$ $\mathcal{LS}(B_2) : \left\{ \begin{array}{l} A = f(A) \\ B = f(B) \\ C = f(C) \\ D = f(D) \end{array} \right \begin{array}{l} 1 = f(A, 2, C) \\ 2 = f(A, B, 3) \\ 3 = f(B, C, 1) \end{array}$ <p>\Rightarrow 4 degrees of freedom</p> <p>\rightarrow stop</p> <p>$\rightarrow B_0$ is not 3-reducible</p>	<p>initial subset:</p> $I(B_0) : \{\underline{2}, 3, A, B, C\}$ $\mathcal{LS}(B_0) : \left\{ \begin{array}{l} A = f(A) \\ B = f(B) \\ C = f(C) \end{array} \right \begin{array}{l} 3 = f(3) \\ 2 = f(A, B, 3) \end{array}$ <p>\Rightarrow 4 degrees of freedom</p> <p>back-tracing: $3 = f(B, C, 1)$</p> $I(B_1) : \{\underline{1}, \underline{2}, \underline{3}, A, B, C\}$ $\mathcal{LS}(B_1) : \left\{ \begin{array}{l} A = f(A) \\ B = f(B) \\ C = f(C) \end{array} \right \begin{array}{l} 1 = f(1) \\ 2 = f(A, B, 3) \\ 3 = f(B, C, 1) \end{array}$ <p>\Rightarrow 4 degrees of freedom</p> <p>back-tracing: $1 = f(A, 2, C)$</p> $I(B_2) : \{\underline{1}, \underline{2}, \underline{3}, A, B, C\}$ $\mathcal{LS}(B_2) : \left\{ \begin{array}{l} A = f(A) \\ B = f(B) \\ C = f(C) \end{array} \right \begin{array}{l} 1 = f(A, 2, C) \\ 2 = f(A, B, 3) \\ 3 = f(B, C, 1) \end{array}$ <p>\Rightarrow 3 degrees of freedom</p> <p>\rightarrow stop</p> <p>$\rightarrow B_0$ is 3-reducible</p>
--	--

Figure 12: Two detailed examples for back-tracing.

in the second example in Figure 12, when back-tracing vertex 1.

3.3 Algorithm to decide 3-reducibility

The algorithm to decide 3-reducibility for a given subset of pseudo-triangles B_0 of a pseudo-triangulation \mathcal{PT} can be designed as a standard graph search on a directed acyclic graph (\mathcal{DAG}).

Algorithm: 3-reducibility

main routine:

```
TEST3RED( $\mathcal{PT}, B_0$ )  
   $I_0 = \{v \mid v \text{ is corner of at least one member of } B_0\}$   
  mark all  $v \in I_0$  that are non-corner of a member of  $B_0$   
  clear  $\mathcal{I}$   
  return SEARCH( $\mathcal{PT}, I_0, 0$ )
```

sub-routines:

```
SEARCH( $\mathcal{PT}, I, n$ )  
  ADD( $I, n, \mathcal{I}$ )  
  if NUMCOMPLETE( $\mathcal{PT}, I$ ) > 3 then  
    return FALSE  
  if NUMUNMARKED( $I$ ) < 4 then  
    return TRUE  
  is3red = FALSE  
  for all  $v \in I$  do  
    if ( not marked( $v, I$ ) ) & (  $v$  incomplete in  $\mathcal{PT}$  ) then  
       $I_{next} = \text{BACKTRACE}(I, v, \mathcal{PT})$   
      if not FIND( $I_{next}, n+1, \mathcal{I}$ ) then  
        is3red = SEARCH( $\mathcal{PT}, I_{next}, n+1$ )  
        if is3red then break  
  return is3red
```

```
BACKTRACE( $I, v, \mathcal{PT}$ )  
   $\{c_1, c_2, c_3\} = \text{GETCORNER}(v, \mathcal{PT})$   
  for all  $c \in \{c_1, c_2, c_3\}$  do  
    if  $c \notin I$  then  
      add  $c$  to  $I$  unmarked  
  mark  $v$  in  $I$   
  return  $I$ 
```

help function legend:

NUMCOMPLETE(\mathcal{PT}, I)...	returns the number of vertices in I that are complete in \mathcal{PT}
NUMUNMARKED(I)	... returns the number of unmarked vertices in I
ADD(I, n, \mathcal{I})	... adds I to the collection \mathcal{I} for the n -th step
FIND(I, n, \mathcal{I})	... returns TRUE if I is found in \mathcal{I} for the n -th step
GETCORNER(v, \mathcal{PT})	... returns the three corners c_1, c_2, c_3 of the unique pseudo-triangle ∇_v in \mathcal{PT} --- equivalent to $v = f(c_1, c_2, c_3)$

This pseudo-code exemplifies a depth-first search algorithm on the \mathcal{DAG} . The graph is directed because each connection represents the addition of a pseudo-triangle to the current set. Only subsets of pseudo-triangulations on the same level of the \mathcal{DAG} can be identical, as for identical subsets the same number of pseudo-triangles has to be added to B_0 . Since, for the same reason, no connections to nodes on a previous level or the same level are possible, the graph is acyclic.

The root is the initial vertex set I_0 respectively the initial subset B_0 of pseudo-triangles, for which we want to decide 3-reducibility. Each node contains a vertex set with less than 3 vertices that are complete in \mathcal{PT} . Further each node has more than 3 *unmarked* vertices. A leaf contains a vertex set with either more than 3 complete vertices or only 3 *marked* vertices. The latter one then is the last visited leaf because in that case the algorithm finishes with " B_0 is 3-reducible".

Remarks A pseudo-triangulation subset B_0 that is 3-reducible has to reside in the same plane within all possible surfaces. If B_0 consists of a single pseudo-triangle then B_0 is obviously 3-reducible. Later, we will specify less trivial initial subsets to gain more advantage of this observation, namely concerning the initial task we want to fulfill — to decide the combinatorial projectivity of \mathcal{PT} .

4 Two Implications of 3-Reducibility

In the last chapter we have defined a powerful tool, the 3-reducibility attribute and its decision algorithm. This chapter shows how to use this knowledge to decide the combinatorial projectivity of a given pseudo-triangulation \mathcal{PT} . It will be advantageous to find small subsets of pseudo-triangles of \mathcal{PT} which are 3-reducible, rather than deciding 3-reducibility for all possible subsets. These subsets have to be well defined, however.

4.1 Sets with 3 height defining vertices

Consider subsets B_0 and B_R of pseudo-triangles of \mathcal{PT} such that B_R can be obtained from repeatedly back-tracing B_0 . We say that B_R *witnesses* the 3-reducibility of B_0 if there exist only three height defining vertices, c_1, c_2, c_3 , for B_R . In our examples so far, each such B_R was enclosed by an induced pseudo-triangle of \mathcal{PT} that had c_1, c_2, c_3 as corners. It is worth to prove whether this is coincidence or not, because it would allow to narrow the range of problem classes down to induced pseudo-triangles of \mathcal{PT} .

Lemma 11 *Let B_0 be a 3-reducible subset of pseudo-triangles of \mathcal{PT} . Let B_R be a subset (with only 3 height defining vertices, c_1, c_2, c_3) that witnesses the 3-reducibility of B_0 . Then B_R forms a connected point set which resides within $\text{conv}(\{c_1, c_2, c_3\})$. Moreover, $\mathcal{F}(B_R)$ lies in the plane spanned by $\{c_1, c_2, c_3\}$ and their heights.*

Proof. The subset B_R was reached by successively adding the unique pseudo-triangles ∇_v to the actual subset B , for vertices v that are height defining for B and incomplete in \mathcal{PT} . Each connected component of pseudo-triangles within B_R has at least 3 height defining vertices that are not height defining for other connected components. Thus there exist $k \cdot 3$ height defining vertices for B_R if B_R contains k connected components. As B_R has only 3 height defining vertices, B_R forms a single connected component.

Now we reverse the back-tracing sequence and insert the incomplete vertex v_n that was back-traced the last. The height of v_n depends on the heights of c_1, c_2, c_3 . Therefore the vertex has to reside on the same plane, by Theorem 1 and within $\text{conv}(\{c_1, c_2, c_3\})$, by Observation 8. The next vertex, v_{n-1} , depends on c_1, c_2, c_3 and v_n . Clearly, this vertex lies in the same plane again. Now there exist 4 possible convex hulls, within v_{n-1} must reside, namely:

$$\text{conv}(\{c_1, c_2, c_3\}), \text{conv}(\{v_n, c_2, c_3\}), \text{conv}(\{c_1, v_n, c_3\}), \text{conv}(\{c_1, c_2, v_n\}).$$

Since v_n lies within $\text{conv}(\{c_1, c_2, c_3\})$, also the three additional convex hulls lie inside $\text{conv}(\{c_1, c_2, c_3\})$, and therefore also v_{n-1} . This proves the lemma by induction. \square

Corollary 1 *Consider a subset B_R as in Lemma 11. Then B_R is combinatorial coplanar.*

Theorem 3 *Let B_0 be an arbitrary subset of pseudo-triangles of \mathcal{PT} . Then B_0 is combinatorial coplanar iff B_0 is 3-reducible.*

Proof. Let B_R witness the 3-reducibility of B_0 . By Corollary 1, B_R is combinatorial coplanar. As $B_0 \subset B_R$, also B_0 is combinatorial coplanar.

Now assume that B_0 is not 3-reducible. This means that there exist at least 4 height defining vertices for B_0 , and also for every subset, B_i , of \mathcal{PT} that can be obtained from repeatedly back-tracing B_0 . Let B_L be the (unique) largest subset obtained from B_0 where every vertex complete in B_L is also complete in \mathcal{PT} . The heights of the at least 4 height defining vertices for B_0 are linearly dependent on at least 4 vertices complete in \mathcal{PT} . Furthermore, the corresponding linear equations are independent, as there would exist a 3-reducible B_i , otherwise. Therefore \mathbf{h} and \mathcal{PT} can always be chosen, such that at least 4 height defining vertices for B_0 are not coplanar. \square

Lemma 12 *Consider B_R and $\{c_1, c_2, c_3\}$ as in Lemma 11. Then the boundary of the union of the pseudo-triangles of B_R forms an induced pseudo-triangle, ∇ , of \mathcal{PT} with corners c_1, c_2, c_3 (and possible holes). Each non-corner of ∇ is incomplete in \mathcal{PT} .*

Proof. Let M be the union of the pseudo-triangles of B_R . By Lemma 11, M resides within $\text{conv}(\{c_1, c_2, c_3\})$, and M forms a connected point set. Therefore, the vertices c_i and c_j of $\{c_1, c_2, c_3\}$ are connected by polygonal chains Z_{ij} that lie on the boundary of M .

Each vertex v on Z_{ij} (except c_i and c_j) is incomplete in B_R (otherwise it would be height defining for B_R) and thus incomplete in \mathcal{PT} , by Observation 3. Additionally, the unique pseudo-triangle where v is a non-corner has to be a member of B_R , otherwise v would be complete in B_R . Recall that Z_{ij} lies on the boundary of M and this implies that Z_{ij} is concave at v . We conclude that Z_{ij} forms a side chain between the two corners c_i and c_j of a pseudo-triangle. \square

4.2 Characterizing combinatorial projectivity

By Theorem 3 we know that if we can find a 3-reducible subset B_0 of pseudo-triangles of \mathcal{PT} , we can conclude that there exist combinatorial coplanar pseudo-triangles in \mathcal{PT} . But we cannot decide the combinatorial projectivity of \mathcal{PT} , in general; see Figure 13.

Combinatorial coplanarity arises iff a subset of at least two pseudo-triangles of \mathcal{PT} can be back-traced to a subset that has only 3 height defining vertices, even if the former pseudo-triangles have no common edge. Examples are illustrated in Figure 13, and for the special case of double adjacency, in Figure 10 (on page 21).

Note that the existence of combinatorial coplanar pseudo-triangles in \mathcal{PT} does not necessarily imply that \mathcal{PT} is combinatorial non-projective. But we will see




Figure 13: Though the gray pseudo-triangles are combinatorial coplanar, both pseudo-triangulations are combinatorial projective.

that if no combinatorial coplanarity exists within \mathcal{PT} then \mathcal{PT} is combinatorial projective.

Fortunately, we know from Observation 7 that a pseudo-triangulation, \mathcal{PT} , in a polygonal region, R , is combinatorial non-projective if \mathcal{PT} contains combinatorial planar edges. By Definition 2.33, each edge of R is combinatorial non-planar. Thus we will concentrate on internal edges of \mathcal{PT} and their neighborhood. If an edge e is combinatorial planar, then two combinatorial coplanar facets have e in common. This means that the two facets and the edge e reside in a plane that is defined by only three vertices. So the subset B_0 consisting of the two pseudo-triangles of \mathcal{PT} adjacent at e can be repeatedly traced back to a subset with only three height defining vertices. That is, B_0 is 3-reducible.

Definition 4.1 (3-Reducible Edge) *If two pseudo-triangles ∇_1 and ∇_2 are adjacent at an edge e and $B_0 = \{\nabla_1, \nabla_2\}$ is 3-reducible, we call e a 3-reducible edge.*

This is **not** a new definition for 3-reducibility. It rather is an abbreviation for a construction rule for well defined pseudo-triangulation subsets:

$$\text{"3-reducible edge } e\text{"} \iff \left\{ \begin{array}{l} \text{"The subset of two pseudo-triangles} \\ \text{adjacent at } e \text{ is 3-reducible."} \end{array} \right.$$

Recall Figure 12 on page 26. In both examples, the initial subset consisted of two adjacent pseudo-triangles. Therefore we can say Figure 12(a) illustrates checking whether $\overline{12}$ is a 3-reducible edge and Figure 12(b) does the same for edge $\overline{A2}$.

We can see from the examples that the combinatorial planar edge $\overline{A2}$ is 3-reducible, whereas the combinatorial non-planar edge $\overline{12}$ is not.

Lemma 13 *An edge e of \mathcal{PT} is combinatorial planar if and only if e is 3-reducible.*

Proof. This is a direct consequence of Theorem 3 and Definition 4.1. \square

We now have a tool ready that allows us to characterize the combinatorial projectivity of \mathcal{PT} .

Theorem 4 *A pseudo-triangulation \mathcal{PT} is combinatorial non-projective if at least one 3-reducible edge exists in \mathcal{PT} . If no 3-reducible edge exists then \mathcal{PT} is combinatorial projective.*

Proof. Assume that a 3-reducible edge e exists within \mathcal{PT} . By Lemma 13, edge e then is combinatorial planar, and by Observation 7 this implies that \mathcal{PT} is combinatorial non-projective.

Now assume that no 3-reducible edge exists within \mathcal{PT} . In other words, all edges of \mathcal{PT} are combinatorial non-planar by Lemma 13. By Definition 2.33 this means that for each edge of \mathcal{PT} the two incident pseudo-triangles are combinatorial non-coplanar.

We still have to argue that there exists a geometrical realization for each combinatorial non-planar edge in a surface, without forcing another edge to be planar. Remember that it is allowed to perturb the vertex set, see Definition 2.30. In addition observe that forcing planarity of an edge e , using a special height vector, means that e is planar only for an exact combination of heights of at least four vertices. By perturbing the height vector with an arbitrarily small ϵ the edge e can be made non-planar without making any other edge planar. \square

Corollary 2 *Each 3-reducible edge of \mathcal{PT} is enclosed by an induced pseudo-triangle of \mathcal{PT} that has exclusively incomplete vertices as non-corners.*

Proof. Assume the pseudo-triangles ∇_1 and ∇_2 are adjacent at the 3-reducible edge e , and let the subset $B_R(e)$ witness the 3-reducibility of the subset $B_0(e) = \{\nabla_1, \nabla_2\}$. By Lemma 12, $B_R(e)$ is enclosed by an induced pseudo-triangle, ∇ , of \mathcal{PT} that has exclusively incomplete non-corners. As $B_0(e) \subset B_R(e)$ holds, also $B_0(e)$ and thus e are enclosed by ∇ . Though this applies only to all internal edges, it is sufficient, because 3-reducible edges have two incident pseudo-triangles by Definition 4.1. \square

Corollary 3 *Edges spanned by two complete vertices are never 3-reducible.*

Proof. Assume that the edge e , spanned by two complete vertices, c_1, c_2 , is 3-reducible, and the subset $B_R(e)$ witnesses the 3-reducibility of the subset $B_0(e) = \{\nabla_1, \nabla_2\}$. By Corollary 2, 3-reducible edges always reside in the interior of an induced pseudo-triangle, ∇ , of \mathcal{PT} . The corners of ∇ plus c_1 and c_2 must not exceed the number of three height defining vertices for $B_R(e)$. As complete vertices are height defining for each subset of pseudo-triangles that contains them, only one additional height defining vertex is allowed for $B_R(e)$ beside c_1 and c_2 . It is easy to see that only one of c_1 and c_2 can be a corner of ∇ . Therefore there exist at least four height defining vertices for $B_R(e)$, which is a contradiction to the assumed 3-reducibility of $B_0(e)$. \square

Observation 10 *In the interior of an induced pseudo-triangle, ∇ , of \mathcal{PT} that encloses exclusively incomplete vertices, each edge is 3-reducible.*

Proof. Let B_R be the set of pseudo-triangles within ∇ . As ∇ encloses only incomplete vertices, ∇ is pointed pseudo-triangulated. So there exist exactly 3 height defining vertices for B_R , namely the corners of ∇ . Thus B_R witnesses the 3-reducibility of $B_0(e) = \{\nabla_1(e), \nabla_2(e)\}$ for each edge e inside ∇ , with $\nabla_1(e), \nabla_2(e)$ being the two pseudo-triangles adjacent at e . \square

Observation 11 *If all induced pseudo-triangles of a pseudo-triangulation, \mathcal{PT} , are empty (for example "real" pseudo-triangles of \mathcal{PT}) then \mathcal{PT} is combinatorial projective.*

5 Combinatorial Stability

Though the concept of 3-reducibility characterizes combinatorial non-projectivity, it is not very manageable. The existing stability definition (Definition 2.35) provides a simpler attribute for pseudo-triangulations that allows to directly conclude to their combinatorial projectivity. Unfortunately, it does not cover all instances of combinatorial non-projectivities. Therefore we need to extend it, if we want to use it further on. Only looking at sets of incomplete vertices has been found insufficient to decide projectivity (see Figure 14), and the additionally needed complete vertices have to be carefully chosen.

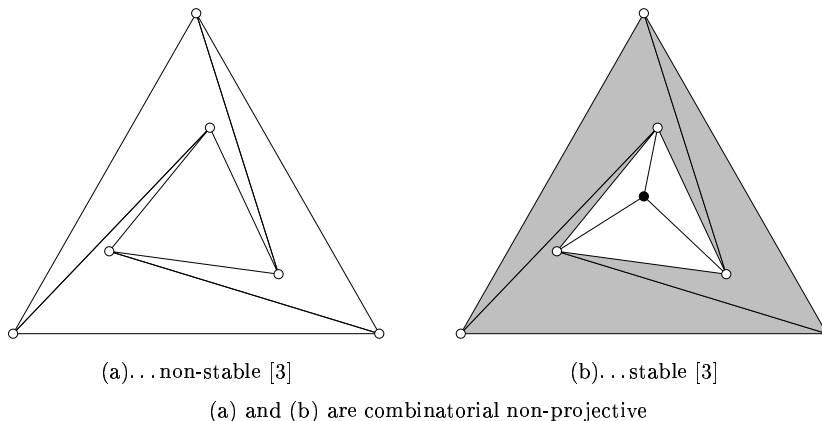


Figure 14: Both examples show combinatorial non-projective pseudo-triangulations. Example (a) from [3] is non-stable whereas (b) is stable. Therefore (b) is a counter-example for deciding combinatorial projectivity with the stability supposed in [3] (Definition 2.35).

Figure 14(b) illustrates a counter-example for the stability attribute supposed in [3] (see, Definition 2.35). No incomplete vertex can be removed without changing the pointedness of the complete vertex (black dot). Thus this pseudo-triangulation would be called stable, but the conclusion to projectivity is incorrect. Compare this with Figure 14(a), where the problematic complete vertex is missing. Here it is possible to remove the set of all three incomplete vertices, maintaining a valid pseudo-triangulation and the pointedness of remaining vertices.

What we will get later on, is a new stability condition that lets us conclude to the existence of combinatorial coplanar facets and further on to the combinatorial projectivity of pseudo-triangulations. With this goal in mind, we have to take a closer look on complete vertices within pseudo-triangles and the different geometrical configurations in which they can appear.

5.1 Finding the correct complete vertices

Consider a pseudo-triangulation \mathcal{PT} with edge set E . From Lemma 12 we know that we can restrict ourselves to finding 3-reducibilities within induced pseudo-triangles,

∇ , of \mathcal{PT} , see Definition 2.18. Further all non-corners of ∇ have to be incomplete in \mathcal{PT} .

Now let ∇ contain a 3-reducible subset B_0 of pseudo-triangles. Assume that the corners of ∇ are the 3 height defining vertices for the subset that witnesses the 3-reducibility of B_0 . This assumption makes sense as otherwise there would exist another induced pseudo-triangle inside ∇ that fulfills the requirements. As complete vertices may exist within ∇ , see Figure 14(b), they have to be "isolated" somehow. This means that repeatedly back-tracing, starting with subset B_0 , never may add a pseudo-triangle that is incident to such complete vertices.

Observe that a vertex v inside ∇ that is complete in \mathcal{PT} is always surrounded by some polygonal cycle, for example by the side chains of ∇ . Analyzing these polygonal cycles we will see that they form a criterion that suffices to decide whether v is "isolated" or not.

Definition 5.1 (Encircled) Consider a pseudo-triangulation \mathcal{PT} with edge set E . Let ∇ be an induced pseudo-triangle of \mathcal{PT} . Further, let \mathcal{IP} be a simple polygon inside ∇ , using only edges of E . A vertex v is called encircled by \mathcal{IP} if v lies in the interior of \mathcal{IP} and \mathcal{IP} does not use any edge of ∇ .

Definition 5.2 (Convex Encircled) Consider ∇ and \mathcal{IP} as in Definition 5.1. A vertex v is called convex encircled by \mathcal{IP} if v is encircled by \mathcal{IP} and \mathcal{IP} is convex.

Definition 5.3 (With Incompletes Encircled) Consider ∇ and \mathcal{IP} as in Definition 5.1. A vertex v is called to be with incompletes encircled by \mathcal{IP} if v is encircled by \mathcal{IP} and each vertex of \mathcal{IP} is incomplete in \mathcal{PT} .




Figure 15: A complete vertex, v , within a pseudo-triangle. Observe that v is convex encircled with incompletes.

Later we will state a definition for a special encirculating polygon, and we will see that asking for convex encirculation is sufficient. But in the meantime we have to look at all possible kinds of encirculations of complete vertices in ∇ .

Our intention is to decide, whether the area inside ∇ but outside each polygonal cycle that "isolates" some complete vertex, contains 3-reducible subsets of pseudo-triangles.

Figure 15 shows a complete vertex, v (black dot), that is with incompletes convex encircled by an internal polygon (bold edges). The pseudo-triangles of $\nabla \setminus \mathcal{IP}$ (gray) are combinatorial coplanar because the encircled complete vertex cannot "influence" the heights of vertices outside its convex "prison".

Lemma 14 *Let B_R be a subset of pseudo-triangles of a pseudo-triangulation \mathcal{PT} , which has exactly 3 height defining vertices. Further, let the polygonal region M be the union of the members of B_R . Consider the induced pseudo-triangle, ∇ , of \mathcal{PT} that arises as outer boundary of M . Then each complete vertex of \mathcal{PT} in the interior of ∇ is convex encircled with incompletes.*

Proof. Assume that there exists some complete vertex, v , inside ∇ (the lemma is trivial, otherwise). The subset B_R does not contain v as a vertex, because there would exist ≥ 4 height defining vertices for B_R , otherwise (the corners of ∇ plus v). Thus the polygonal region M has a hole, H , with v inside H .

If H has a reflex vertex, r , then r is a corner of M , and therefore r is a complete vertex in B_R . So r is a fourth height defining vertex for the 3-reducibility witness B_R — a contradiction. We conclude that H is a valid (by Definition 5.1) convex encirculation of v . Further, the hole H is with incompletes convex encircling v , because B_R has only 3 height defining vertices, and complete vertices of \mathcal{PT} are always height defining. \square




Figure 16: The complete vertex, v , in example (a) is with incompletes non-convex encircled. Example (b) derives from (a) by flipping an edge. Here v is convex encircled with incompletes.

Figure 16(a) illustrates an exemplification for complete vertices that are with incompletes non-convex encircled. The complete vertex v (black dot) is only non-convex encircled with incompletes by the inner polygon \mathcal{IP} (bold). The arrows indicate the path of height dependence. The reflex vertex r on \mathcal{IP} can be back-traced to v in one step, as v is corner of the unique pseudo-triangle where r is non-corner. Further, r itself is corner of another two vertices' unique pseudo-triangle, and so on.

Figure 16(b) derives from Figure 16(a) by flipping one edge inside the non-convex incomplete \mathcal{IP} (bold dashed). Note that now v is convex encircled with incompletes (bold) and thus "isolated". Therefore the subset of pseudo-triangles (gray) between ∇ and the (bold) convex polygon is 3-reducible.

Definition 5.4 (Component (of Polygons)) *Let S be a vertex set and let \mathcal{P} be a set of polygons with vertices in S . Then $P_1, P_2 \in \mathcal{P}$ belong to the same component of \mathcal{P} if they share at least 2 vertices of S . If X is a component of polygons, containing the polygons $P_1 \dots P_k$ then we consider X as the polygonal region $\bigcup_{i=1}^k P_i$.*

Lemma 15 *Let ∇ be an induced pseudo-triangle of \mathcal{PT} all whose non-corners are incomplete in \mathcal{PT} . Assume that each complete vertex, v , of \mathcal{PT} interior to ∇ is convex encircled by some polygon $H(v)$. Then the subset, B , of pseudo-triangles that forms the polygonal region $\nabla \setminus \mathcal{H}$, has 3 height defining vertices, with $\mathcal{H} = \bigcup H(v)$, for v complete in \mathcal{PT} .*

Proof. We claim that each component of polygons, X , of \mathcal{H} is a convex polygon. If X equals some encirculation $H(v)$ then X is convex by assumption. Otherwise, X is the non-disjoint union of at least two convex encirculations. Assume the existence of a reflex vertex, r , of X . Then r is a vertex of at least two convex encirculations. Therefore no two consecutive edges incident to r span an angle $> \pi$. Thus r is a complete vertex of \mathcal{PT} that is not encircled — a contradiction.

We show next that each vertex of X is incomplete in the subset, B , of pseudo-triangles forming $\nabla \setminus \mathcal{H}$. As X is convex, each vertex w of X has its angle $> \pi$ inside $\nabla \setminus \mathcal{H}$. As all complete vertices interior to ∇ are convex encircled by assumption, w has to be incomplete in \mathcal{PT} . This implies that w is incomplete in B .

Finally, each non-corner, k , of ∇ is incomplete in B as well, because k is incomplete in \mathcal{PT} by assumption. In summary, $\nabla \setminus \mathcal{H}$ has exactly 3 vertices being complete in (respectively, height defining for) B , namely ∇ 's corners. \square

Remarks Note that each component of polygons X of \mathcal{H} is convex with incompletes encircling some vertices that are complete in \mathcal{PT} .

Figure 17 shows an exemplification to the proof of Lemma 15. As these are only simplified examples, not all edges of the underlying pseudo-triangulation are shown. Black dots represent complete vertices, white dots incomplete ones. The solid, dashed and dash-dotted polygons indicate the different convex encirculations of complete vertices.

If each encirculation of a complete vertex within some induced pseudo-triangle is convex then each component of the union of all such convex encirculations is convex. Further, each vertex of these components is incomplete.

Figure 17(a) illustrates an example where the two convex encirculations for v_1 and v_2 join to one convex component. Figure 17(b) exemplifies the contradiction construction where the component formed by the two convex encirculations of v_1 and v_2 is not convex and therefore results (in this example) in two unencircled




Figure 17: Two complete vertices (black dots), are complete encircling each other and building a component of polygons.

complete vertices (r_1, r_2) . Figure 17(c) exemplifies how the (correct) component for Figure 17(b) could look like.

Corollary 4 *Consider ∇ as in Lemma 15. Each complete vertex inside ∇ is with incompletes convex encircled iff each such vertex is convex encircled.*

Proof. Assume each complete vertex v inside ∇ is convex encircled by some $H(v)$. By Lemma 15, the subset B of pseudo-triangles that forms $\nabla \setminus \bigcup H(v)$ has only 3 height defining vertices. As the 3 corners of ∇ are height defining for B , no additional vertex can be complete in B . Therefore each hole of the region $\nabla \setminus \bigcup H(v)$ has to be convex with exclusively incomplete vertices on its boundary.

On the other hand, it is trivial that each with incompletes convex encircled vertex is also convex encircled. \square

Definition 5.5 (Smallest Convex Encirculation) *Let S be a vertex set with some pseudo-triangulation and let $v \in S$ be some complete vertex. Let $\mathcal{P}(v)$ be the set of all convex polygons spanned by S that encircle v . Further we define for a polygon P the function $d(P)$ that counts the number of vertices of S that lie in the interior of P . Then $P_s(v) \in \mathcal{P}(v)$ is the smallest convex encirculation of v if $d(P_s(v)) < d(P_i(v))$, for each $P_i(v) \in (\mathcal{P}(v) \setminus P_s(v))$.*

Definition 5.6 (Convex Inner Polygon, \mathcal{CIP}) *Consider ∇ as in Lemma 15. Let $H(v)$ be the smallest convex encirculation of a vertex v that is complete in \mathcal{PT} and that lies inside ∇ . Let $\mathcal{H} = \bigcup H(v)$, for each vertex v inside ∇ that is complete in \mathcal{PT} . Then each component of \mathcal{H} is called a convex inner polygon (of ∇).*

Compare this definition for \mathcal{CIP} 's with Lemma 15. Instead of an arbitrary convex encircling polygon for each v , as in Lemma 15, let's take the \mathcal{CIP} for each v . Thereby we obtain the following corollary:

Corollary 5 *Consider ∇ as in Lemma 15. Let each complete vertex, v , inside ∇ be convex encircled. Then the subset, B , of pseudo-triangles that forms the polygonal region $\nabla \setminus \bigcup \mathcal{CIP}$ has exactly 3 height defining vertices, with $\bigcup \mathcal{CIP}$ being the union of all convex inner polygons of ∇ .*




Figure 18: Using edges of ∇ for forming an enclosing polygonal cycle (bold) is not allowed.

Finally we want to argue, why using edges of the enclosing polygon, ∇ , is forbidden for encirculations in Definition 5.1. Figure 18 shows three different examples of what happens, when edges of ∇ were used for encirculations of complete vertices. Using some edge, like the bold edge in Figure 18(a), for a convex "hole", as indicated with the dotted edges, is not allowed, because the non-corners of ∇ (square vertices) must not become complete. Thus using such an edge leads to a non-convex "hole", see Figure 18(b), which does exclude the existence of 3-reducible subsets of pseudo-triangles inside $\nabla \setminus \bigcup \mathcal{CIP}$ (see Lemmata 14 and 15 and Definition 5.6 resp. Corollary 5).

Therefore the only way using an edge of ∇ for a convex "hole", partitions the problem into two subproblems, see Figure 18(c), that can be handled separately. Firstly, the convex "hole" itself. Secondly, the remaining, gray, induced pseudo-triangle. An intuitive reason for prohibiting edges of ∇ is that this will not lead to "holes" inside ∇ . It rather leads to cutting off a piece of ∇ .




Figure 19: Using only vertices but no edge of ∇ for forming an enclosing polygonal cycle (bold) is allowed.

But remember, using only vertices of ∇ is allowed for an encirculation of a complete vertex. Compare Figure 19 with the three examples in Figure 18.

5.2 The updated stability property

Using the previously acquired knowledge, we are able to define a substructure of pseudo-triangulations that contains 3-reducible subsets.

Definition 5.7 (Vicious Pseudo-Triangle) *Let \mathcal{PT} be a pseudo-triangulation. An induced pseudo-triangle, ∇ , of \mathcal{PT} is called vicious if (1) there exists at least*

one vertex in the interior of ∇ , (2) all non-corners of ∇ are incomplete and (3) all complete vertices internal to ∇ are convex encircled.

Remarks The first requirement ensures that ∇ is not a real (empty) pseudo-triangle of \mathcal{PT} . The second and the third requirement could be combined to a single one: $\nabla \setminus \bigcup \mathcal{CIP}$ has to restrict \mathcal{PT} to a pointed pseudo-triangulation. $\bigcup \mathcal{CIP}$ is the union of all \mathcal{CIP} 's in the interior of ∇ . Observe that $\nabla \setminus \bigcup \mathcal{CIP}$ forms an induced punched pseudo-triangle (see Definition 6.2 on page 43).

The vicious pseudo-triangle is a structure that emerges directly from the previous sections. In Section 6 we will investigate this structures more intensive and state a definition for maximal vicious pseudo-triangles called *maximal punched sets* (Definition 6.1).

Definition 5.8 (Combinatorial Stability) *A pseudo-triangulation \mathcal{PT} is called combinatorial stable if no vicious pseudo-triangle of \mathcal{PT} exists. \mathcal{PT} is called combinatorial non-stable, otherwise.*

Remarks This includes (the correct part of) the former stable definition, Definition 2.35, from [3]. If the interior of ∇ (see Definition 5.7) consists of exclusively incomplete vertices, they can be easily eliminated (along with their incident edges). What remains is a valid pseudo-triangulation, since only objects in the interior of ∇ have been eliminated. Also the pointedness of all other vertices remains unchanged, as all non-corners of ∇ have been pointed before.

The (new) combinatorial stability can be decided by considering only the structure of the pseudo-triangulation. We do not have to know the surface for the pseudo-triangulation. It is not possible to directly conclude from combinatorial stability to combinatorial projectivity. But soon we will be able to decide combinatorial projectivity using combinatorial stability and another simple condition that also derives from the pseudo-triangulation. However, first we show that combinatorial stability enables us to directly conclude to combinatorial coplanarity.

Theorem 5 *A pseudo-triangulation, \mathcal{PT} , contains a set of at least two combinatorial coplanar pseudo-triangles iff \mathcal{PT} is combinatorial non-stable.*

Proof. Assume that \mathcal{PT} is combinatorial non-stable. By Definition 5.8, \mathcal{PT} then induces a pseudo-triangle, ∇ , that satisfies all conditions stated in Lemma 15. As ∇ is non-empty, there exists a subset, B , of pseudo-triangles of ∇ , with only 3 height defining vertices and $|B| \geq 2$. So, by Theorem 3, B is combinatorial coplanar.

Now assume that \mathcal{PT} contains a combinatorial coplanar subset, B_0 , of at least two pseudo-triangles. It follows that B_0 is 3-reducible, by Theorem 3. Moreover, by Lemma 12, the outer boundary, the pseudo-triangle ∇ , of the witness B_R of B_0 has only non-corners being incomplete in \mathcal{PT} . As B_R lies inside ∇ , ∇ is non-empty. Finally, by Lemma 14, each complete vertex of \mathcal{PT} inside ∇ is convex encircled. We conclude that ∇ proves \mathcal{PT} to be combinatorial non-stable. \square

Lemma 16 *If a pseudo-triangulation \mathcal{PT} is combinatorial stable, then no 3-reducible edge exists within \mathcal{PT} .*

Proof. By Theorem 5, a combinatorial stable pseudo-triangulation, \mathcal{PT} , contains no combinatorial coplanar pseudo-triangles. Thus no 3-reducible subset of pseudo-triangles of \mathcal{PT} exists by Theorem 3, and this also excludes the existence of 3-reducible edges, see Definition 4.1. \square

Lemma 17 *Let ∇ be an induced pseudo-triangle of \mathcal{PT} that witnesses the combinatorial non-stability of \mathcal{PT} . Then each edge, e , inside ∇ that is not part of any convex inner polygon, is a 3-reducible edge of \mathcal{PT} .*

Proof. Let \mathcal{H} be the union of all \mathcal{CIP} 's inside ∇ . By Corollary 5, each subset of pseudo-triangles of the polygonal region $\nabla \setminus \mathcal{H}$ is 3-reducible. Since e has to be interior to $\nabla \setminus \mathcal{H}$, e is incident to a 3-reducible set of two pseudo-triangles. Therefore e is a 3-reducible edge, see Definition 4.1. \square

Observation 12 *Double-adjacencies, also the deformed type, are ruled out in a combinatorial stable pseudo-triangulation, \mathcal{PT} .*




Figure 20: Double adjacencies are ruled out in combinatorial stable pseudo-triangulations. (a) shows a classical double adjacency with two common edges. (b) illustrates a deformed double adjacency with only one common edge.

Proof. Assume that the two pseudo-triangles ∇_1 and ∇_2 are in (deformed) double-adjacency (sharing no, one, or two edges). As double-adjacent pseudo-triangles have 3 vertices in common, ∇_1 and ∇_2 are combinatorial coplanar by Observation 5. By Theorem 5, \mathcal{PT} is combinatorial non-stable if it contains combinatorial coplanarities. \square

See Figure 20 for an illustration of Observation 12. Incomplete vertices are white dots, complete vertices are black. The 3 vertices that are shared by the

double-adjacent pseudo-triangles are labeled as a , b , and c . The dotted edges are combinatorial planar respectively 3-reducible. Note that the union (bold) of the two double-adjacent pseudo-triangles forms an induced pseudo-triangle that has exclusively incomplete non-corners.

Theorem 6 (Projectivity Theorem) *Let \mathcal{PT} be a pseudo-triangulation. If \mathcal{PT} is combinatorial stable, then \mathcal{PT} is combinatorial projective. If \mathcal{PT} is combinatorial non-stable, then let $\{\nabla_1 \dots \nabla_k\}$ be the set of all induced pseudo-triangles of \mathcal{PT} that witness its combinatorial non-stability. Then \mathcal{PT} is combinatorial non-projective if and only if there exists an edge e of \mathcal{PT} inside some $\nabla \in \{\nabla_1 \dots \nabla_k\}$ that lies in the exterior of all \mathcal{CTP} 's.*

Proof. Using the previously acquired knowledge the proof follows almost immediately. By Theorem 4, \mathcal{PT} is combinatorial projective iff no 3-reducible edge exists. If an edge e of \mathcal{PT} inside some $\nabla \in \{\nabla_1 \dots \nabla_k\}$ exists, then, by Lemma 17, e is 3-reducible. Finally, Lemma 16 excludes the existence of 3-reducible edges if \mathcal{PT} is combinatorial stable. \square

Finally, we want to update an observation from [3] that was formulated with their stable-condition (Definition 2.35).

Observation 13 *Every pointed pseudo-triangulation (with at least two pseudo-triangles) in a polygonal region with only 3 corners is combinatorial coplanar.*

6 Appearance of Hidden Edges

So far we have provided an attribute for a pseudo-triangulation to decide its (combinatorial) projectivity, namely the combinatorial stability. Now we will analyze the class of pseudo-triangulations where the former stability definition (Definition 2.35 from [3]) was insufficient. We will see that pseudo-triangulations may arise which are not only combinatorial non-projective after an improving surface flip but also stay combinatorial non-projective after the following trivial flip. This leaves a pseudo-triangulation whose surface does not project back to a pseudo-triangulation any more.

What we get, is a surface containing planar pseudo-triangles as facets that are punched with convex non-planar holes, the so-called *CIP*'s from Definition 5.6. We define:

Definition 6.1 (Maximal Punched Set) *Let \mathcal{PT} be a pseudo-triangulation. Let M be the polygonal region $\nabla \setminus \bigcup \mathcal{CIP}$, with $\bigcup \mathcal{CIP}$ being the union of all convex inner polygons of a given induced pseudo-triangle ∇ of \mathcal{PT} . Then M is called a maximal punched set if the restriction of \mathcal{PT} to M , $\mathcal{PT}|_M$, is a maximal set of pseudo-triangles with only three height defining vertices.*

Compared to Definition 5.7 (page 39) a maximal punched set is a maximal vicious pseudo-triangle. Observe that a maximal punched set is a maximal combinatorial coplanar face set. The surface of such a face set projects back to a single face that may contain holes and therefore fails to be a pseudo-triangle. Thus we want to define a new class that is a relaxation of pseudo-triangulations and allows pseudo-triangles with internal convex holes.

Definition 6.2 (Punched Pseudo-Triangle) *A polygonal region with exactly three corners is called a punched pseudo-triangle.*

Definition 6.3 (Punched Pseudo-Triangulation) *Let R be a polygonal region. A punched pseudo-triangulation, $\mathcal{PT}_p(R)$, in R , is a cell complex in R whose cells are punched pseudo-triangles.*

Comparing the definitions for punched pseudo-triangulations, pseudo-triangulations (Definition 2.13) and triangulations (Definition 2.7) we see that triangulations and pseudo-triangulations are also punched pseudo-triangulations, because triangles and pseudo-triangles are punched pseudo-triangles. But observe that a punched pseudo-triangle also permits interior convex holes.

Observation 14 *A maximal punched set, M , of a pseudo-triangulation, \mathcal{PT} , is an induced punched pseudo-triangle of \mathcal{PT} . The faces of $\mathcal{PT}|_M$ define a pointed pseudo-triangulation.*

A maximal punched set that consists of only one pseudo-triangle is obviously combinatorial coplanar. By comparing maximal punched sets that consist of more than one pseudo-triangle, with the definition for combinatorial stability (Definition 5.8), we can make the following observation:

Observation 15 *Each maximal punched set, M , of a pseudo-triangulation \mathcal{PT} witnesses the combinatorial non-stability of \mathcal{PT} , provided that $\mathcal{PT}|_M$ contains at least two pseudo-triangles.*

Remarks Note that, alternatively to Definition 5.8, we could say that a pseudo-triangulation, \mathcal{PT} , is combinatorial stable if no induced punched pseudo-triangle, ∇ , with at least one internal vertex exists where $\mathcal{PT}|_{\nabla}$ defines a pointed pseudo-triangulation.




Figure 21: Hidden edges can reappear after improving surface flips.

Definition 6.4 (Hidden Edges) *Consider a pseudo-triangulation, \mathcal{PT} , and a maximal punched set, M , thereof. An edge, e , of \mathcal{PT} is called a hidden edge if e is an inner tangent (Definition 2.24) of M .*

Hidden edges of a pseudo-triangulation, \mathcal{PT} , remain after a trivial flip but are missing in the projection of the surface of \mathcal{PT} . Nevertheless, hidden edges are necessary to maintain a valid pseudo-triangulation.

Observe that hidden edges reappear if a flip leads back to a combinatorial projective pseudo-triangulation. See Figure 21 for an example. Figure 21(a) shows a starting triangulation where flipping (edge-removing resp. planarizing) the two bold edges leads to a punched pseudo-triangulation, shown in Figure 21(b). Only if this punched pseudo-triangulation contains the hidden (dashed) edge it is also a valid pseudo-triangulation. Flipping the bold edge in Figure 21(b) leads to a (valid) pseudo-triangulation where the former hidden edge is not hidden any more.

Observation 16 *Hidden edges are combinatorial planar.*

Proof. By Definition 6.4, each hidden edge, e , is an inner tangent of a maximal punched set, M , of some pseudo-triangulation, \mathcal{PT} . It is easy to see that a single pseudo-triangle prohibits inner tangents, thus $\mathcal{PT}|_M$ has to contain at least two pseudo-triangles. By Observation 15, M witnesses the combinatorial non-stability of \mathcal{PT} and by Lemma 17, e is 3-reducible. Finally, Lemma 13 proves that e is combinatorial planar. \square

Remarks Combinatorial planar, respectively 3-reducible, edges are not necessarily hidden. For example, think of two pseudo-triangles in double adjacency. They are partitioned by two edges that are combinatorial planar but not hidden. Only combinatorial planar edges that remain after a trivial-flip are hidden. In other words: Combinatorial planar edges between combinatorial non-planar vertices are hidden.

Observation 17 *A pseudo-triangulation, \mathcal{PT} , is combinatorial non-projective if at least one hidden edge exists in \mathcal{PT} .*

Proof. By Observation 16 and Lemma 13 each hidden edge is 3-reducible. The combinatorial non-projectivity of \mathcal{PT} follows by Theorem 4. \square

Recall that inner tangents (Definition 2.24) of maximal punched sets are possible hidden edges. As we know the number of inner tangents within some maximal punched set from the number of disjoint internal convex holes, we can provide an upper bound on the number of hidden edges used in a pseudo-triangulation compared to the total number of convex inner polygons.

Lemma 18 *There exist at most $3 \times |\mathcal{CIP}|$ hidden edges, where $|\mathcal{CIP}|$ is the number of convex inner polygons of the pseudo-triangulation.*

Proof. First recall that hidden edges are inner tangents (by Definition 6.4) of some maximal punched set. If all \mathcal{CIP} 's are disjoint, $3 \times |\mathcal{CIP}|$ inner tangents are needed to pseudo-triangulate the maximal punched sets in a pointed way, by Lemma 7.

The number of hidden edges is reduced if some \mathcal{CIP} 's are touching (directly connected to) other \mathcal{CIP} 's and/or the outer boundary of the maximal punched set they form. In this way it is possible to replace all inner tangents (resp. hidden edges) with \mathcal{CIP} 's, see Figure 13 on page 31. \square

Generally there may also exist additional combinatorial planar edges within some maximal punched set that are incident to at least one combinatorial planar vertex (see Definition 2.34). A trivial flip removes all combinatorial planar vertices along with (some of) their incident edges. Therefore we can say, though a trivial flip does not always delete all combinatorial planar edges, it at least minimizes their number.

6.1 The deformation

Definition 6.5 (Deformation, $(M, \mathcal{PT}|_M)$) A pseudo-triangulation, \mathcal{PT} , contains a deformation, $(M, \mathcal{PT}|_M)$, if there exists a maximal punched set, M , of \mathcal{PT} that admits at least one inner tangent. Two deformations, $(M_1, \mathcal{PT}|_{M_1})$ and $(M_2, \mathcal{PT}|_{M_2})$, of \mathcal{PT} are equivalent if $M_1 = M_2$.

Remarks A deformation is characterized by the maximal punched set M and the restriction of \mathcal{PT} to M . By Observation 14, $\mathcal{PT}|_M$ defines a pointed pseudo-triangulation. Two maximal punched sets (see Definition 6.1), $M_1 = \nabla_1 \setminus (\bigcup \mathcal{CIP})_1$ of \mathcal{PT}_1 and $M_2 = \nabla_2 \setminus (\bigcup \mathcal{CIP})_2$ of \mathcal{PT}_2 , are equal, if $\nabla_1 = \nabla_2$ and $(\bigcup \mathcal{CIP})_1 = (\bigcup \mathcal{CIP})_2$.

Observation 18 A pseudo-triangulation that contains a deformation is combinatorial non-projective.




Figure 22: Example for a simple deformation. A maximal punched set (bold) with three hidden edges (dashed).

If a pseudo-triangulation, \mathcal{PT} , contains a deformation, the projection of its surface $\mathcal{F}(\mathcal{PT})$ is a cell complex that is no pseudo-triangulation any more, but a punched pseudo-triangulation. See Figure 22 for an example. The projection of the surface of this pseudo-triangulation does not yield the hidden edges (dashed), and the combinatorial coplanar pseudo-triangles (gray) form a punched pseudo-triangle.

A deformation may evolve by applying a sequence of improving surface flips (and trivial flips) to an initial surface for a given (pseudo-)triangulation. As our surface flips are defined for pseudo-triangulations only, this is an undesirable effect. To reach the maximal locally convex surface, a continuous sequence of improving surface flips has to be guaranteed.

Figure 23 displays the three different possible cases that may occur in one improving surface flip. The starting (pseudo-)triangulation \mathcal{PT} is assumed to be combinatorial projective. In case (a), an improving surface flip results in another combinatorial projective pseudo-triangulation, \mathcal{PT}' . Thus the following trivial flip does nothing (respectively is not necessary). The improving surface flip in case (b)

\mathcal{PT}	$\xrightarrow{\text{ISF}}$	\mathcal{PT}'	$\xrightarrow{\text{TF}}$	\mathcal{PT}''
(a)		comb. projective		comb. projective
(b)		comb. non-projective		comb. projective
(c)		comb. non-projective (\exists deformation)		comb. non-projective (\exists hidden edge)
ISF ... improving surface flip			TF ... trivial flip	

Figure 23: The three cases of an improving surface flip.

leads to a combinatorial non-projective pseudo-triangulation, \mathcal{PT}' . After the following trivial flip, the final pseudo-triangulation \mathcal{PT}'' is again combinatorial projective. Observe that the conclusion from the stability definition from [3] (Definition 2.35) to combinatorial projectivity is correct for the cases (a) and (b).

The last case (c), is exactly the case where this stability definition fails to be the correct condition for combinatorial projectivity. Now the improving surface flip leads to a pseudo-triangulation \mathcal{PT}' that is not only combinatorial non-projective but also contains a deformation. Therefore, the following trivial flip can only minimize the number of combinatorial planar edges, but is not able to completely eliminate them. The consequence is that within the resulting pseudo-triangulation \mathcal{PT}'' so-called hidden edges (Definition 6.4) remain. And this leaves also \mathcal{PT}'' combinatorial non-projective.




Figure 24: Example of the 2 phases of a deformation creation.

Note that the existence of hidden edges implicates the existence of deformations, whereas the existence of deformations does not necessarily imply the existence of hidden edges. Figure 24 exemplifies this instance. Figure 24(a) shows a pseudo-triangulation, where a deformation is prevented by the complete vertex v . Flipping edge e we get the pseudo-triangulation, \mathcal{PT}' (compare with Figure 23), shown in Figure 24(b), where v is now incomplete. The maximal punched set, displayed with bold edges, restricts \mathcal{PT}' to a pointed pseudo-triangulation, whose faces (gray) are combinatorial coplanar. Observe that there exists no single hidden edge in \mathcal{PT}' . Figure 24(c) shows the pseudo-triangulation, \mathcal{PT}'' , after the trivial flip that deletes all combinatorial planar vertices (including v) and thereby minimizes the number of combinatorial planar edges. The remaining combinatorial planar edges are hidden

edges (dashed) that pointed pseudo-triangulate the maximal punched set (bold). Note that there exist 6 possible hidden edges per \mathcal{CIP} , by Lemma 5. Figure 24(c) shows only one possible combination of hidden edges.

Examining the relationship of deformations with \mathcal{CIP} 's, we can state another observation:

Observation 19 *Let \mathcal{PT} be a pseudo-triangulation. If there exists no \mathcal{CIP} in \mathcal{PT} , then \mathcal{PT} contains no deformation, $(M, \mathcal{PT}|_M)$.*

Proof. By Definition 6.5, M must contain at least one inner tangent. If there exists no \mathcal{CIP} , then each maximal punched set M is a pseudo-triangle. But there exists no inner tangent in the interior of a pseudo-triangle. \square

The previously mentioned hidden edges (Definition 6.4) provide an ancillary tool to handle the undesired case, when deformations arise. Using hidden edges, the punched pseudo-triangulation turns into a pseudo-triangulation, and thus the improving surface flips are again applicable.

6.2 Leaving deformations

To find a way out of a deformation, $(M, \mathcal{PT}|_M)$, it is sufficient to look at the convex inner polygons, $\bigcup \mathcal{CIP}$, that form the maximal punched set, $M = \nabla \setminus \bigcup \mathcal{CIP}$, along with the induced pseudo-triangle, ∇ . Nevertheless, we will also take a quick look at ∇ and also at the edges that form the pointed pseudo-triangulation, $\mathcal{PT}|_M$, to show that concentrating on the \mathcal{CIP} 's is not only sufficient, but also the only reliable way to leave a deformation.

Keep in mind that we now operate with a pseudo-triangulation that contains at least one deformation and hidden edges *after* a trivial flip. In Figure 23, this pseudo-triangulation is marked as \mathcal{PT}'' .

Observation 20 *The hidden edges cannot be flipped by improving surface flips.*

Proof. Applying improving surface flips to hidden edges is prohibited by Definition 2.37. \square

Remarks Note that flipping hidden edges (as it is allowed during trivial flips) is only possible by edge-exchanging flips and results in other hidden edges.

Observation 21 *Let $(M, \mathcal{PT}|_M)$ be a deformation in a pseudo-triangulation \mathcal{PT} , where M is a maximal punched set with outer boundary ∇ . Destroying $(M, \mathcal{PT}|_M)$ by flipping a reflex edge of ∇ is only possible by applying an edge-exchanging flip and may result in other deformations.*

Proof. The requirement of reflex edges is obvious for improving flips (Definition 2.37). It is also easy to see that an edge-removing flip will only lead to a bigger




Figure 25: An edge-exchanging flip on the outer boundary of a maximal punched set of a deformation may result in creating other deformations.

outer boundary, since edge-removing flips merge two pseudo-triangles to one. Further, as an edge-exchanging flip changes ∇ , it destroys M and thereby the actual deformation. But, as exemplified in Figure 25, it may happen that this creates other deformations. \square

We showed that destroying a deformation by flipping hidden edges is not allowed, and that flipping edges of the outer boundary of a maximal punched set forming the deformation is not always possible. Even if it is possible, it is not guaranteed to result in a deformation-free pseudo-triangulation (even if the destroyed deformation was the only one). In Figure 25 the actual deformation is destroyed by flipping edge e , of the outer boundary of the maximal punched set (shown bold), to edge e' . The new edge e' splits the old induced pseudo-triangle into two new ones, each containing a \mathcal{CIP} and forming an outer boundary of a new maximal set. Thus, two new deformations were created by e' , as deformations are equivalent if their maximal punched sets are the same (Definition 6.5).

As mentioned before, we will now show that it is always possible to destroy the \mathcal{CIP} 's in a pseudo-triangulation with deformations, thereby also destroying the deformations themselves. We accomplish this in two steps: At first we show that it is possible to destroy a single \mathcal{CIP} by application of improving surface flips (along with trivial flips). Next we prove that we can destroy all \mathcal{CIP} 's with a finite sequence of flips, even though it is possible that destroying a \mathcal{CIP} results in other ones.

Lemma 19 *It is always possible to destroy a convex inner polygon within a deformation by application of (a finite number of) improving surface flips (and trivial flips).*

Proof. Assume that \mathcal{PT} is a pseudo-triangulation in some polygonal region R . Further let \mathcal{CIP} be a convex inner polygon within a deformation in \mathcal{PT} . We prove that \mathcal{CIP} can be destroyed by application of a finite number of improving and trivial surface flips. There exist two cases:

(1) Suppose that at least one edge of \mathcal{CIP} is reflex. Then this edge can be flipped which destroys the \mathcal{CIP} .

(2) Now assume that all edges of \mathcal{CIP} are convex: As each vertex of \mathcal{CIP} has at least one incident reflex edge (see Lemma 9) and the exterior of \mathcal{CIP} is combinatorial coplanar, the reflex edges exist in the interior of \mathcal{CIP} . As all vertices of \mathcal{CIP} are incomplete in $\mathcal{PT}|_{R \setminus \mathcal{CIP}}$, flipping in the interior of \mathcal{CIP} does not change the heights of those vertices. On the other hand, all vertices of \mathcal{CIP} are complete in $\mathcal{PT}|_{\mathcal{CIP}}$, because \mathcal{CIP} is convex. So flipping in the interior of \mathcal{CIP} cannot change their heights in *any* surface defined by \mathcal{CIP} alone. Thus improving surface flips in the interior of \mathcal{CIP} can convexify the part of the surface for \mathcal{PT} that lies inside \mathcal{CIP} , $F|_{\mathcal{CIP}}$, by the Optimality Theorem (Theorem 2). Convexifying the interior of \mathcal{CIP} leads to one of two possible results:

(a) The convexified surface $F|_{\mathcal{CIP}}$ is planar: In this case a trivial flip removes all vertices interior to \mathcal{CIP} , in particular, all previously complete vertices. Thereby \mathcal{CIP} is destroyed.

(b) The convexified surface $F|_{\mathcal{CIP}}$ is strictly convex: Now \mathcal{CIP} has to have reflex edges and can be destroyed as in case (1).

Finally, the number of flips used is finite by Theorem 2. \square

Theorem 7 *Let \mathcal{PT} be a pseudo-triangulation containing deformations. It is always possible to find a finite sequence of improving surface flips and trivial flips for \mathcal{PT} that leads back to a deformation-free pseudo-triangulation, \mathcal{PT}' , using hidden edges for maintaining pseudo-triangular cell complexes during flipping.*

Proof. Let $(M, \mathcal{PT}|_M)$ be a fixed deformation of \mathcal{PT} . Let ∇ be the outer boundary of M . We show: improving flipping that destroys $(M, \mathcal{PT}|_M)$ is possible, without changing \mathcal{PT} in the exterior of ∇ .

By Lemma 19, it is possible to destroy a single convex inner polygon with a finite number of improving surface flips (including trivial flips).

Doing so it is possible to create additional \mathcal{CIP} 's in the (former) interior of the destroyed \mathcal{CIP} . In this case, the additional \mathcal{CIP} 's encircle at least one complete vertex less than the destroyed \mathcal{CIP} did, because a \mathcal{CIP} is a smallest convex encirculation (Definition 5.5). Thus either at least one complete vertex is now unencircled or at least one complete vertex turned incomplete (or was removed). In the first case, the deformation that contained the destroyed \mathcal{CIP} , has been destroyed, by Definition 6.1 and Definition 6.5. In the second case, we can delete the additionally created \mathcal{CIP} 's with a finite sequence of flips (Lemma 19). As we "lose" at least one complete vertex for each destroyed \mathcal{CIP} , only a finite number of additional \mathcal{CIP} 's can be created and thus the number of applied improving surface flips and trivial flips stays finite.

The creation of an additional \mathcal{CIP} is also possible in the (former) exterior of the destroyed \mathcal{CIP} . In this case, the new and "bigger" \mathcal{CIP} may encircle each complete vertex that was encircled by the destroyed \mathcal{CIP} . Fortunately, this new \mathcal{CIP} can only be created by using hidden edges as edges of the new \mathcal{CIP} and/or by merging with other \mathcal{CIP} 's. When merging with other \mathcal{CIP} 's, the number of \mathcal{CIP} 's

is reduced by at least 1, as no other \mathcal{CIP} 's emerge in exchange. When using hidden edges to save the \mathcal{CIP} , some hidden edges become edges of the boundary of this \mathcal{CIP} and therefore are not hidden any more. As no additional hidden edges arise this way, the number of hidden edges is also reduced by at least 1. Thus in both cases the number of flips again is finite, as both hidden edges and \mathcal{CIP} 's only exist in a finite number. \square

6.3 Bypassing deformations

We just found a way to handle deformations of pseudo-triangulations using hidden edges. But hidden edges are only ancillary edges that do not really exist in the projection of a surface. Thus using hidden edges is no elegant approach. It would be better if we will not have to use them, thereby never reaching a pseudo-triangulation containing deformations. If we could show that we can always avoid creating deformations, we could guarantee flipping sequences to the optimal surface using only improving and trivial surface flips.

To accomplish this, it is sufficient to show that we can find a flipping sequence to avoid a particular deformation we were right about to create. To find a way around a deformation, we have to first identify the edges that create the deformation and then find other reflex edges beside the undesired ones.

Suppose flipping the reflex edge e results in a deformation and in the case of an edge-exchanging flip in the edge e' . Of what types are e and e' ?

Lemma 20 *Flipping e to e' creates a deformation $(M, \mathcal{PT}|_M)$ if e' is an edge of M or the removal of e leads to a pointed pseudo-triangulation $\mathcal{PT}|_M$.*




Figure 26: The different edge categories for creating a deformation.

Proof. Since the deformation did not exist before the flip, either the new edge e' has to be part of it or the flipped edge e was preventing it. The partitioning into

these cases (see Figure 26 for examples) follows directly from Definition 6.5 (and Definition 6.1). \square

Figure 26 shows examples of flips that can create a deformation. Figure 26(a) exemplifies the creation of the maximal punched set M . If the edge e_3 has already been removed this can be achieved by either flipping edge e_1 to e'_1 (if e_2 has already been flipped to e'_2) or by flipping edge e_2 to e'_2 (if e_1 has already been flipped to e'_1). The edge e_3 keeps the vertex v_1 complete, thus an edge-removing flip of e_3 would also create a deformation (if edges e_1 and e_2 have already been flipped to e'_1 and e'_2). Figure 26(b) shows another version for creating a deformation with an edge-removing flip. Flipping either edge e_{4a} or edge e_{4b} turns the vertex v_2 from complete to incomplete, and thus the restriction of the pseudo-triangulation to the maximal punched set is a pointed pseudo-triangulation.

Lemma 21 *Suppose that flipping edge e creates a deformation, $(M, \mathcal{PT}|_M)$. Then not flipping at most three edges, namely e and at most two other edges, forever prevents $(M, \mathcal{PT}|_M)$.*




Figure 27: Three edges keep a vertex complete.

Proof. By Lemma 20, we can divide the proof into two parts:

(1) Assume e flips to edge e' of M . If e' crosses e , then M is obviously prevented forever if e will not be flipped. If e' has one endpoint, v , in common with e , then v is non-corner of M . The simultaneous existence of e and e' makes v non-pointed by Lemma 2 and therefore $\mathcal{PT}|_M$ can never be a pointed pseudo-triangulation. By Definition 6.5, this prevents $(M, \mathcal{PT}|_M)$ forever. Observe, that the common vertex v cannot be a corner of M , because in this case the deformation already would have existed before.

(2) Assume e was preventing $\mathcal{PT}|_M$ from being pointed. This means, the edge-removing flip applied to e makes one non-pointed vertex, v , pointed. The vertex v is either a vertex of M or in the interior of M . If v is a vertex of M , then v is kept non-pointed by either e alone or a potential flipping pair (Definition 2.23) e and e' . Therefore not flipping e and possibly e' prevents the deformation forever.

If v lies in the interior of M , then exactly three edges, e , e' and e'' , suffice to keep v non-pointed, see Figure 27 for an example. The edges e' and e'' are the

previous respectively next edge of e incident to v . This means that there lie no additional edges between e' and e respectively between e and e'' but there may exist an arbitrary number of edges between e'' and e' . The angle between e' and e'' , within which the edge e lies, has to be greater than π , because v would not be pointed without e , otherwise. Thus the opposite angle between e' and e is smaller than π . Furthermore, also the angles between e and e' , and e and e'' have to be smaller than π , as v is non-pointed with e . Therefore not flipping any of the three edges, e , e' and e'' , prevents $(M, \mathcal{PT}|_M)$ forever. \square

Lemma 22 *If flipping edge e creates a deformation, there must exist other reflex edges beside e that are flipable without creating any deformation.*

Proof. We assume that the deformation $(M, \mathcal{PT}|_M)$ and the associated \mathcal{CIP} 's arise after flipping edge e . By Lemma 20, the edge e either flips to an edge e' that is an edge of M or e is removed by the flip and $\mathcal{PT}|_M$ becomes a pointed pseudo-triangulation. As M consists of an induced pseudo-triangle ∇ and various \mathcal{CIP} 's (see Definition 6.1) we get three different cases: (1) the construction of ∇ , (2) making $\mathcal{PT}|_M$ a pointed pseudo-triangulation and (3) the construction of a \mathcal{CIP} .

Case (1) and (2): In both cases each \mathcal{CIP} interior to ∇ already exists. We claim that at least one \mathcal{CIP} has a reflex edge on its boundary or in its interior, before the flip. Assume the contrary. Then, for each \mathcal{CIP} , the surface above \mathcal{CIP} and its adjacent faces of \mathcal{PT} is convex. On the other hand, as a deformation arises after the flip, the \mathcal{CIP} 's were either surrounded by a pseudo-quadrilateral (case (1)) or ∇ already existed with an additional 4th height defining vertex in its interior (case (2)). Either way, the height of a single additional height defining vertex cannot force all edges of \mathcal{CIP} to be convex.

So there exists a reflex edge (a) in the interior of \mathcal{CIP} or (b) on the boundary of \mathcal{CIP} . In case (a), we try to flip this edge. If this causes no new deformation (in the interior of \mathcal{CIP}) we are done. Otherwise, we solve the problem recursively. In case (b), flipping the boundary edge of \mathcal{CIP} is save: All vertices inside \mathcal{CIP} are complete, as incomplete vertices have to be incident to reflex edges, and case (a) would have occurred. Flipping an edge of \mathcal{CIP} destroys it and at least one complete vertex within the former \mathcal{CIP} becomes an additional height defining vertex within ∇ .

Case (3): The deformation is constructed by closing a \mathcal{CIP} . A \mathcal{CIP} has at least 3 vertices. Each of them is incomplete, and at least one vertex, v , has to lie in the interior of ∇ and is not incident to edge e .

Flipping any reflex edge incident to v (which has to exist by Lemma 9) is either save or it causes a new deformation. But this cannot happen by closing another \mathcal{CIP} , as there exists at least one additional height defining vertex within ∇ (more specifically, within the \mathcal{CIP} to-be). Therefore such a new deformation can only be created if the flip closes a new induced pseudo-triangle within ∇ (that contains at least one \mathcal{CIP}). But this subproblem can be solved recursively as in case (1). \square

Corollary 6 *If the edge e is the only remaining reflex edge, it is not possible that flipping e will create a deformation.*

Proof. This follows directly from Lemma 22. There always exist safely flipable reflex edges, if flipping e would create a deformation. Therefore a deformation cannot arise, if e is the only reflex edge left. \square

Theorem 8 *It is always possible to find a sequence of improving surface flips (including the trivial flip), to reach the optimal surface without ever creating a deformation.*

Proof. This is an immediate consequence of Lemma 22. \square

7 Summary

7.1 Application of theory

As mentioned in the introduction (page 1), from preliminary work of this thesis a computer program to calculate and represent a surface from a pseudo-triangulation exists. Further this program allows flipping by clicking to reach the optimal surface (see Definition 2.42). As it was not clear how to treat planar edges of the surface, the program did not flip to the optimum with an arbitrary flipping sequence.

Using the newly achieved insight we were able to update the program. We adapted the trivial flip (see Definition 2.40) with respect to the new theory and used the hidden edges (see Definition 6.4) to maintain pseudo-triangular cell complexes. Using the updated program we constructed some examples of surfaces that are displayed in the appendix (page 57ff).

In Appendix A we will exemplify a complete sequence of "flipping to optimality", in Figures 28.31. We start with a triangulation (first row of images in Figure 28) and apply improving surface flips (Definition 2.37), until the optimal surface is reached (see Theorem 2 on this). A step in the flipping sequence is made up of an improving surface flip followed by a trivial flip. There are three columns for each step; column (a) displays the 2D-view (the pseudo-triangulation) and the columns (b) and (c) display the surface in 3D-view from different angles. After 15 flips the maximal locally convex surface is reached (last row in Figure 31).

Appendix B exemplifies another surface. During the represented sequence of improving flips, there emerge deformations (Definition 6.5) in this surface. As in Appendix A we show three different views of the surface per step. But this time we do not display the complete sequence (each step). We only show the most important respectively interesting steps. The initial triangulation is shown in the first row of Figure 32. After 6 improving surface flips the first deformation arises, see the second row of Figure 32. Another ("smaller") deformation arises after the 9th flip, shown in the last row of Figure 32. The first two rows in Figure 33 show the destruction of the "smaller" deformation after the 12th flip, and then the surface after the 16th flip. Finally the last two rows in Figure 33 represent the last deformation after the 21st flip and the optimal surface that has been reached after a total of 24 improving surface flips.

7.2 Conclusion

The goal of this work was to "repair" the stability condition for pseudo-triangulations from [3]. We investigated the mechanisms of coplanarity and introduced a new concept, the *3-reducibility*. We demonstrated how to use the 3-reducibility to detect combinatorial coplanar face sets. We described the interrelationship between 3-reducibility and *combinatorial projectivity* and stated a new condition for pseudo-triangulations, the *combinatorial stability*. Using 3-reducibility we were able to

prove that the combinatorial projectivity of a pseudo-triangulation can be decided with the new combinatorial stability condition.

Further, we exemplified that, starting with arbitrary triangular surfaces, improvident flipping may lead to surfaces which do not project to pseudo-triangular cell complexes. According to this circumstance we proved that there always exist sequences of improving (and trivial) flips that avoid leaving the class of pseudo-triangulations when flipping to optimality.

7.3 Future Work

We introduced the new concept of *combinatorial stability* for pseudo-triangulations. Nevertheless, it remains to find an efficient algorithm to decide combinatorial stability respectively combinatorial projectivity of pseudo-triangulations.

In [3] flipping algorithms are provided to prove a quadratic upper bound for the length of flipping sequences to optimality for convex underlying domains as well as vertex empty simple polygons as underlying domain. Those algorithms might be adapted to bypass deformations. Moreover, the effect of the existence of punched pseudo-triangulations on the flip distance to optimality should be investigated. Also the flip distance to optimality of arbitrary polygonal regions is still unknown, see [3].

In this thesis, a new type of cell complex, the punched pseudo-triangulation, has been introduced. Punched pseudo-triangulations are relaxations of pseudo-triangulations, but in contrast to pseudo-triangulations, flipping is not defined for punched pseudo-triangulations up to now. When executing an edge-exchanging flip, as defined for pseudo-triangulations, within two punched pseudo-triangles, the new edge, defined by a geodesic, may cross holes of the punched pseudo-triangles. Thereby a flip within two punched pseudo-triangles may result into several new (punched) pseudo-triangles, depending on the number of holes crossed by the geodesic. We plan to elaborate on how to perform this flip in a unique way.

Punched pseudo-triangles are polygonal regions with exactly three corners. The possibility of holes in such regions has certain relations to the fact that pseudo-simplices in three-space may contain tunnels, see [6]. We believe that a thorough study of pseudo-triangulations will shed additional light into the class of "triangulation-relaxing" cell complexes.

A Flip Sequence to Optimum




Figure 28: Improving surface flip sequence to optimum, 01/04.




Figure 29: Improving surface flip sequence to optimum, 02/04.




Figure 30: Improving surface flip sequence to optimum, 03/04.




Figure 31: Improving surface flip sequence to optimum, 04/04.

B Surface Examples with Deformation

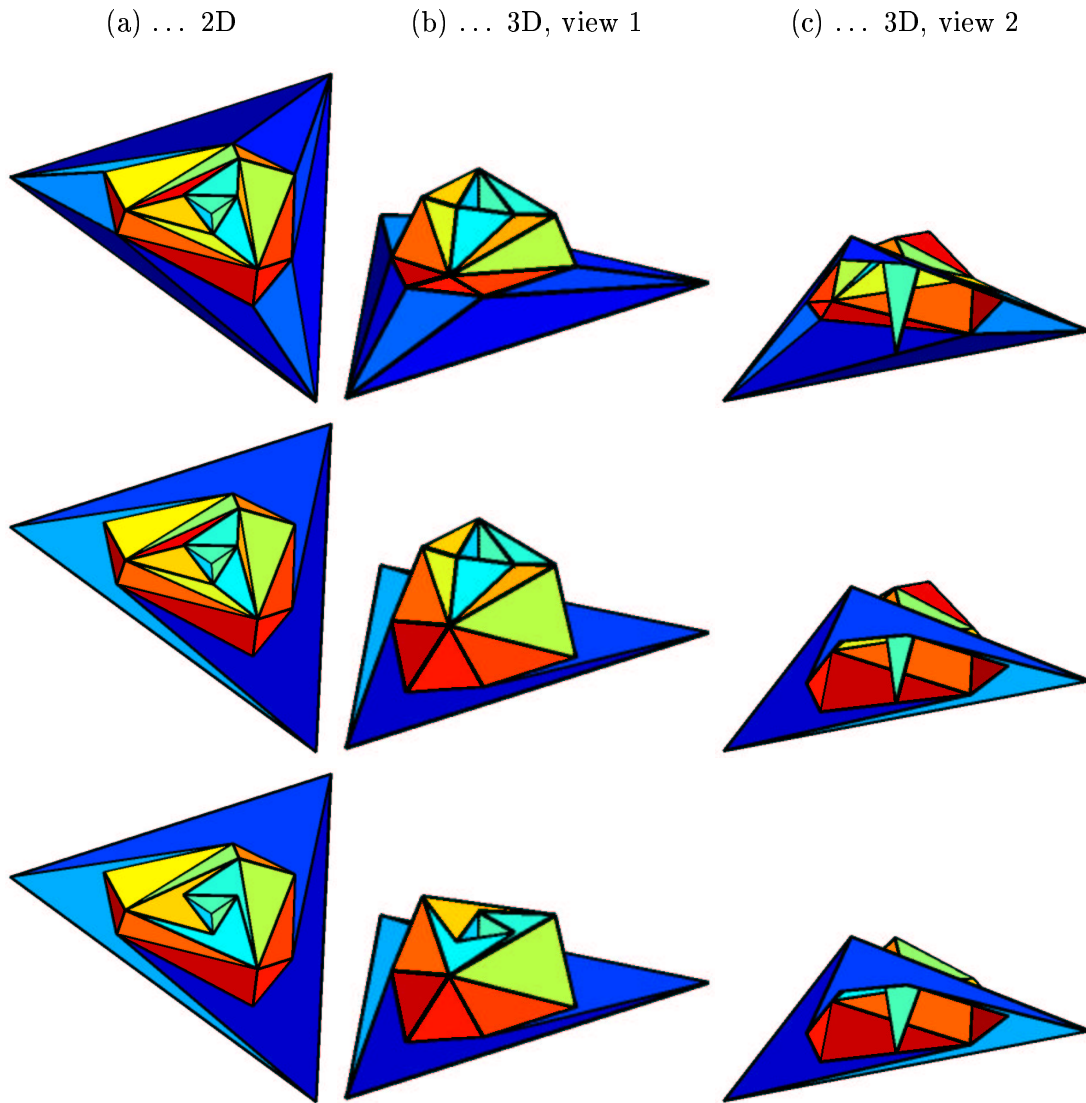


Figure 32: Example surface with deformation during flipping to optimum, 01/02.

(a) ... 2D

(b) ... 3D, view 1

(c) ... 3D, view 2




Figure 33: Example surface with deformation during flipping to optimum, 02/02.

References

- [1] P.K.Agarwal, J.Basch, L.J.Guibas, J.Hershberger, and L.Zhang, *Deformable free space tilings for kinetic collision detection*. In Algorithmic and Computational Robotics: New Directions (Proceedings of the 5th Workshop on the Algorithmic Foundations of Robotics), B.R.Donald, K.Lynch, and D.Rus, editors, 2001, 83-96.
- [2] O.Aichholzer, F.Aurenhammer, F.Hurtado, *Sequences of spanning trees and a fixed tree theorem*. Computational Geometry 21 (2002), 3-20.
- [3] O.Aichholzer, F.Aurenhammer, H.Krasser, P.Brass, *Pseudo-triangulations from surfaces and a novel type of edge flip*. SIAM Journal on Computing, 32:1621-1653, 2003.
- [4] O.Aichholzer, F.Aurenhammer, H.Krasser, *Enumerating order types for small point sets with applications*. Order, 19:265-281, 2002.
- [5] O.Aichholzer, H.Krasser, *The point set order type database: A collection of applications and results*. 13th Canadian Conference on Computational Geometry (CCCG), Waterloo, Ontario, Canada, pages 17-21, 2001.
- [6] F.Aurenhammer, H.Krasser, *Pseudo-simplicial complexes from maximal locally convex functions*. Japan Conference on Discrete & Computational Geometry (JCDCG 2004), Tokyo, Japan, 2004, to be presented.
- [7] J.-D.Boissonnat, M.Yvinec, *Algorithmic geometry*. Cambridge University Press, Cambridge, UK, 1998.
- [8] H.Borouchaki, P.-L.George, *Delaunay triangulation and meshing: Application to finite elements*. Editions Hermes, Paris, 1998.
- [9] B.Chazelle, H.Edelsbrunner, M.Grigni, L.J.Guibas, J.Hershberger, M.Sharir, and J.Snoeyink, *Ray shooting in polygons using geodesic triangulations*. Algorithmica 12 (1994), 54-68.
- [10] H.Edelsbrunner, N.R.Shah, *Incremental topological flipping works for regular triangulations*. Algorithmica 15 (1996), 223-241.
- [11] J.E.Goodman, R.Pollack, *Multidimensional sorting*. SIAM Journal on Computing, 12:484-507, 1983.
- [12] M.T.Goodrich, R.Tamassia, *Dynamic ray shooting and shortest paths in planar subdivisions via balanced geodesic triangulations*. Journal of Algorithms 23 (1997), 51-73.

- [13] R.Haas, D.Order, G.Rote, F.Santos, B.Servatius, H.Servatius, D.Souvaine, I.Streinu, and W.Whiteley, *Planar minimally rigid graphs and pseudo-triangulations*. Computational Geometry 2003, 154-163.
- [14] D.Kirkpatrick, J.Snoeyink, B.Speckmann, *Kinetic collision detection for simple polygons*. International Journal of Computational Geometry & Applications 12 (2002), 3-27.
- [15] D.Kirkpatrick, B.Speckmann, *Kinetic maintenance of context-sensitive hierarchical representations for disjoint simple polygons*. In Proceedings of the 18th Annual ACM Symposium on Computational Geometry, Barcelona, Spain, 2002, 179-188.
- [16] H.Krasser, *Order types of point sets in the plane*. PhD thesis, Institute for Theoretical Computer Science, Graz University of Technology, Austria, October 2003.
- [17] C.L.Lawson, *Properties of n -dimensional triangulations*. Computer Aided Geometric Design 3 (1986), 231-246.
- [18] M.Pocchiola, G.Vegter, *Topologically sweeping visibility complexes via pseudo-triangulations*. Discrete & Computational Geometry 16 (1996), 419-453.
- [19] M.Pocchiola, G.Vegter, *Minimal tangent visibility graphs*. Computational Geometry 6 (1996), 303-314.
- [20] M.Pocchiola, G.Vegter, *On polygonal covers*. In Advances in Discrete and Computational Geometry, B.Chazelle, J.E.Goodman, and R.Pollack, editors, Contemporary Mathematics 223, AMS, Providence, 1999, 257-268.
- [21] F.P.Preparata, M.I.Shamos, *Computational Geometry: An Introduction*. Springer-Verlag, New York, 1985.
- [22] G.Rote, F.Santos, I.Streinu, *Expansive motions and the polytope of pointed pseudo-triangulations*. In Discrete and Computational Geometry—The Goodman-Pollack Festschrift, B.Aronov, S.Basu, J.Pach, and M.Sharir, editors, Algorithms and Combinatorics, Springer, Berlin, 2003, 699-736.
- [23] B.Speckmann, C.D.Tóth, *Allocating vertex π -guards in simple polygons via pseudo-triangulations*. In Proceedings of the 14th ACM-SIAM Symposium on Discrete Algorithms, Baltimore, MD, 2003, SIAM, Philadelphia, 2003, 109-118.
- [24] I.Streinu, *A combinatorial approach to planar non-colliding robot arm motion planning*. In Proceedings of the 41st IEEE Symposium on Foundations of Computer Science (FOCS), Redondo Beach, CA, 2000, 443-453.