

Generalized swap operation for tetrahedrizations

B. Lehner¹, B. Hamann², G. Umlauf³

¹Department of Computer Science, University of Kaiserslautern, Germany
`lehner@cs.uni-kl.de`

²Institute for Data Analysis and Visualization (IDAV), Department of Computer Science, University of California, Davis, USA
`hamann@cs.ucdavis.edu`

³Department of Computer Science, HTWG Constance, Germany
`umlauf@htwg-konstanz.de`

Abstract. Mesh optimization of 2D and 3D triangulations is used in multiple applications extensively. For example, mesh optimization is crucial in the context of adaptively discretizing geometry, typically representing the geometrical boundary conditions of a numerical simulation, or adaptively discretizing the entire space over which various dependent variables of a numerical simulation must be approximated. Together with operations applied to the vertices the so-called edge or face swap operations are the building block of all optimization approaches. To speed up the optimization or to avoid local minima of the function measuring overall mesh quality these swaps are combined to generalized swap operations with a less local impact on the triangulation.

Despite the fact that these swap operations change only the connectivity of a triangulation, it depends on the geometry of the triangulation whether the generalized swap will generate inconsistently oriented or degenerate simplices. Because these are undesirable for numerical reasons, this paper is concerned with geometric criteria that guarantee the generalized swaps for a 3D triangulation to yield only valid, non-degenerate triangulations.

Key words: 3D triangulation; geometric conditions; swap operations.

1 Introduction

Triangulations of points in 2D space for a mesh of triangles or points in 3D space for a mesh of tetrahedra are crucially important for numerous applications encountered in scientific and engineering application, including numerical simulation, shape approximation, or visualization. In scattered data approximation [8, 15, 20] 2D triangulations are used to define a piecewise linear coarse approximation of a dense data set, assigning a “height value” for every vertex. This technique can also be used for image compression [4, 5, 18, 21] and video compression [17, 19]. For reverse engineering [1, 6, 9, 12], the 2-manifold surface to be reconstructed is approximated by a 3D triangulation that contains no tetrahedra. For mechanical engineering and physical simulations [14, 24], 3D triangulations are used as meshes for finite element methods.

For all of these applications the triangulation needs to be optimized with respect to an application-dependent cost function measuring mesh quality based on a multitude of proper mesh quality variables, including, for example, point distribution, approximation error [7, 18], triangle shape [10], dihedral angles [14], etc. The optimization process is usually based on simple, local changes in the triangulations such as repositioning of vertices [15], insertions and removal of vertices [7, 11] and edge and face swaps [22]. While the first of these operations change geometry and connectivity of the triangulation the swaps change only the connectivity of a triangulation. To speed up the optimization or to avoid local minima during mesh optimization multiple edge and face swaps are combined to generalized swap operations that change the connectivity of more than three tetrahedra of the triangulation [13, 17, 23, 25], see Section 2.

However, it depends on the geometry of the triangulation if a generalized swap will generate flipped or degenerate simplices. We present in this paper geometric criteria that guarantee that a generalized swap operation in a 3D triangulation will generate only valid, non-degenerate triangulations.

2 Related Work

In general, a *swap operation* replaces d -dimensional simplices of a triangulation ($d \geq 1$) by other simplices. It usually affects only a local area of the triangulation, and changes the connectivity of the triangulation without changing the number or position of the vertices.

Lawson [16] was among the first scientists studying and publishing swap operations systematically. He showed that $d + 2$ points in d dimensions, which do not all lie in a hyper-plane, have either one unique triangulation \mathcal{T} or two possible triangulations \mathcal{T}_1 and \mathcal{T}_2 . Which case happens depends on the vertex positions, see Figure 1 for the 2D case. In the latter case, \mathcal{T}_1 and \mathcal{T}_2 differ only in connectivity and the transformation from \mathcal{T}_1 to \mathcal{T}_2 is called swap operation $s_{1 \rightarrow 2}(\mathcal{T}_1) = \mathcal{T}_2$. The opposite transformation is $s_{2 \rightarrow 1}(\mathcal{T}_2) = \mathcal{T}_1$. Because $s_{1 \rightarrow 2} \circ s_{2 \rightarrow 1} = s_{2 \rightarrow 1} \circ s_{1 \rightarrow 2} = \text{id}$, $s_{1 \rightarrow 2}$ and $s_{2 \rightarrow 1}$ are inverse operations.

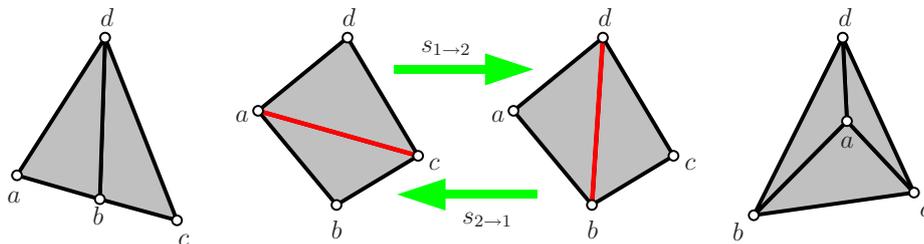


Fig. 1. Triangulations of four points in the 2D case.

If \mathcal{T}_1 is a subset of a larger triangulation \mathcal{T} , the swap operation can be applied by replacing only the simplices of \mathcal{T}_1 with those of \mathcal{T}_2 , and leaving all simplices of \mathcal{T} unchanged, i.e., $\mathcal{T}' = (\mathcal{T} \setminus \mathcal{T}_1) \cup \mathcal{T}_2$. Note that the subset \mathcal{T}_1 has to be a triangulation, i.e. it has to fill the convex hull of its vertices, and must be convex.

Additionally to these basic swaps, one can construct generalized swap operations that replace a set of simplices C of the triangulation by a different set of simplices C' . Thus, C and C' are not required to cover the convex hull of their vertices. Since the generalized swaps are usually more powerful, they can lead to a good triangulation with less swap operations, but are often less efficient.

One way to construct a generalized swap operation is to combine a sequence of basic swap operations to a so-called *composed swap operation*. For the 2D case Yu et al. [25] use a combination of two edge swap operations. If a simple edge swap does not reduce the cost function, they swap the edge and one of its adjacent edges. Thus, the affected faces do not need to form a convex polygon for the composed swap operations. Using the composed swap operations can improve the optimization results significantly.

Concerning the 3D case, the set of swap operations is larger and more varied than in the 2D case. Again, we can categorize them into basic swap operations and composed swap operations. According to Lawson [16], there are five different settings of five points a, b, c, d and e in 3D space, only two of which have two different triangulations and therefore provide swap operations, see Figure 2. If three points are collinear, or four points a, b, c, d are coplanar with $d \in \text{conv}(a, b, c)$, or $e \in \text{conv}(a, b, c, d)$ there is only one possible triangulation, see Figures 2 I., III., and V. If exactly four points are coplanar and form a convex quadrilateral q there are two possible triangulations with flipped diagonals of q , see Figure 2 II. Because the triangulation consists of two cells before and after the swap, the swap is called a *2-2 swap*. For the most general case in which all five points are corners of $\text{conv}(a, b, c, d, e)$ there are also two possible triangulations, see Figure 2 IV. Because this swap replaces three cells by two and vice versa, it is called a *3-2 swap* or *2-3 swap*, respectively.

When applied to a subset of a triangulation \mathcal{T} , the 2-2 swap is only possible if the two faces $\{a, b, d\}$ and $\{b, c, d\}$ are border faces of \mathcal{T} . If they are interior faces, the incident two cells also have to be swapped, see Figure 3. This leads to the 4-4 swap, which replaces four cells with four other cells.

In 3D also a combination of basic swap operations can be more powerful. Joe [13] systematically analyzed the possible settings. Every face of a triangulation is assigned to nine different categories, describing their local setting and their status of being transformable by a basic swap operation. He proposes a set of composed swap operations to transform faces that are initially not transformable, by first swapping adjacent faces. For every composed swap operation, he lists the cells that are removed and created. From this list, he provides criteria in [13] to compute the change of a cost function c resulting from each of the operations, if c is the minimum of the costs of the individual cells.

Another class of composed swap operations is the class defined by the generalizations of the 3-2 and 2-3 swaps, see [3, 24].

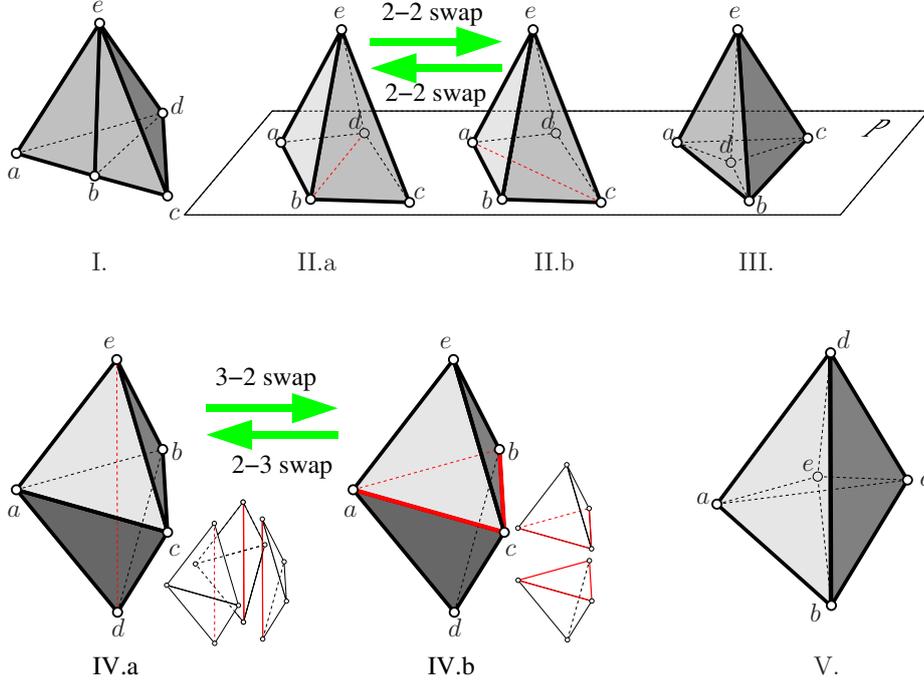


Fig. 2. The different settings of five points in the 3D case.

Generalized 3-2 swap (\mathcal{G}_{32}) A *generalized 3-2 swap* (\mathcal{G}_{32}) can be applied to an edge $e = \{a, b\}$ with $n \geq 3$ incident cells $C = \{c_1, \dots, c_n\}$, with $c_i = \{a, b, v_i, v_{i+1}\}$ and $v_{n+1} \equiv v_1$, see Figure 4 (left). The loop (v_1, \dots, v_n) is split into a set of $n - 2$ connected faces $F = \{f_1, \dots, f_{n-2}\}$. Note that the choice of F is not unique. \mathcal{G}_{32} replaces the edge e with the faces F , where the n cells C are replaced by the $2(n - 2)$ cells $C' = \{c'_{a,1}, c'_{b,1}, \dots, c'_{a,n-2}, c'_{b,n-2}\}$ with $c'_{a,i} = f_i \cup \{a\}$ and $c'_{b,i} = f_i \cup \{b\}$.

Generalized 2-3 swap (\mathcal{G}_{23}) We say a face $f = \{v_1, v_2, v_3\}$ is *sandwiched* between vertices a and b , if the two cells incident to f are $c_1 = \{a, v_1, v_2, v_3\}$ and $c_2 = \{b, v_1, v_2, v_3\}$. A *generalized 2-3 swap* (\mathcal{G}_{23}) is applied to a set $F = \{f_1, \dots, f_{n-2}\}$ of faces, which are sandwiched between two points a and b , see Figure 4 (right). A new edge $e = \{a, b\}$ is inserted into the triangulation, and the border edges of F are connected to the new edge e to form the new cells. Let $C' = \{c'_{a,1}, c'_{b,1}, \dots, c'_{a,n-2}, c'_{b,n-2}\}$ be the set of cells incident to the faces $c'_{a,i} = f_i \cup \{a\}$ and $c'_{b,i} = f_i \cup \{b\}$ of F , and (v_1, \dots, v_n) be the loop of vertices defined by the border edges of F . \mathcal{G}_{23} replaces the faces of F by the edge $e = \sigma_{\{a,b\}}$, and the $2(n - 2)$ cells of C' are replaced by the n cells $C = \{c_1, \dots, c_n\}$, with $c_i = \{a, b, v_i, v_{i+1}\}$, and $v_{n+1} \equiv v_1$.

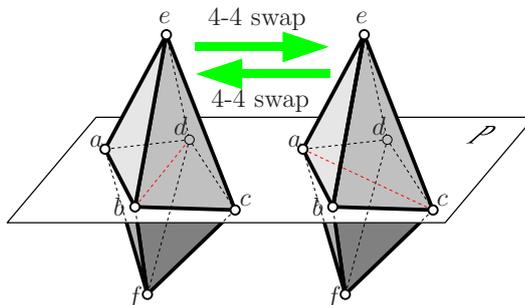


Fig. 3. The 4-4 swap is used if the faces of the 2-2 swap are no border faces.

\mathcal{G}_{23} is the inverse of \mathcal{G}_{32} . Since the choice of faces is not unique in either direction, applying the one swap operation after the other leads to the start triangulation only if for both swaps the same faces are chosen. Also note that the 2-3 swap is a special case of \mathcal{G}_{23} , the 3-2 swap of \mathcal{G}_{32} , and the 4-4 swap a special case of \mathcal{G}_{23} and \mathcal{G}_{32} .

The execution of \mathcal{G}_{32} and \mathcal{G}_{23} can result in invalid triangulations. In Sections 4 and 5 we discuss necessary and sufficient geometric conditions to ensure the validity of the resulting triangulation. Shewchuk [23] notes that these swaps can be replaced by a series of 2-3 and 3-2 swaps, where the intermediate triangulations are topologically correct, but may contain degenerate or inverted cells. In Section 6 we show that there is always a sequence of 2-3, 3-2, and 4-4 swaps to replace a \mathcal{G}_{23} or \mathcal{G}_{32} swap without degenerate or inverted cells.

3 Notation

In order to define the generalized swap operation in terms of connectivity changes and associated geometric conditions, we first adjust our notation properly.

A 3D triangulation $\mathcal{T} = (\mathcal{V}, \mathcal{E}, \mathcal{F}, \mathcal{C})$ (tetrahedrization) consists of a set of vertices \mathcal{V} , edges $\mathcal{E} \subset \mathcal{V}^2$, faces $\mathcal{F} \subset \mathcal{V}^3$ (triangles), and cells $\mathcal{C} \subset \mathcal{V}^4$ (tetrahedra). Thus, an edge is a pair of vertices, a face a triple of vertices, and a cell a quadruple of vertices. All these entities are ordered such that \mathcal{T} is an oriented simplicial 3-complex, where the edges of adjacent faces and the faces of adjacent cells are order reversely. In this case, we call \mathcal{T} a *valid* triangulation. We will use set operations to define new faces and cells, i.e., for $v_1 \in \mathcal{V}$, $e = (v_2, v_3) \in \mathcal{E}$, $f = (v_2, v_3, v_4) \in \mathcal{F}$ and $c \in \mathcal{C}$ we define

$$\begin{aligned} e \cup \{v_1\} &= (v_1, v_2, v_3) \in \mathcal{F}, \\ f \cup \{v_1\} &= (v_1, v_2, v_3, v_4) \in \mathcal{C} \end{aligned}$$

and

$$\begin{aligned} e \in f &\iff (v_2, v_3) \text{ is a sub-tuple of } f, \\ f \in c &\iff (v_2, v_3, v_4) \text{ is a sub-tuple of } c. \end{aligned}$$

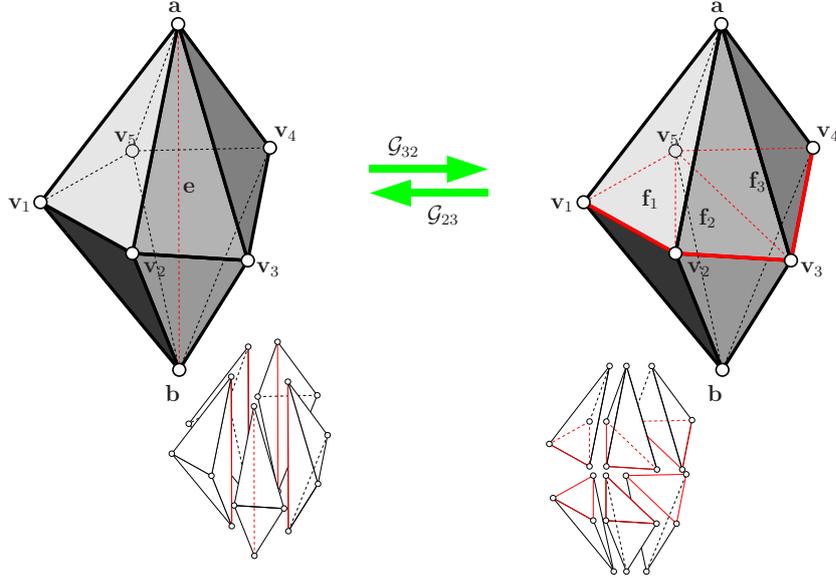


Fig. 4. The generalized 3-2 and 2-3 swaps.

While \mathcal{V} , \mathcal{E} , \mathcal{F} , and \mathcal{C} describe only the connectivity of the triangulation, a geometric realization of \mathcal{T} is defined by associating a point $\mathbf{v} \in \mathbb{R}^3$ to every vertex $v \in \mathcal{V}$. The geometric realizations of an edge $e \in \mathcal{V}^2$, a face $f \in \mathcal{V}^3$, or a cell $c \in \mathcal{V}^4$ are then defined as the convex hull of the geometric realizations of their vertices, and are also denoted in boldface letters \mathbf{e} , \mathbf{f} , and \mathbf{c} , respectively. Furthermore, for a set M of edges, faces, or cells, we denote by \mathbf{M} the union of the geometric realizations of the elements of M . Throughout this paper, geometric realizations of elements of a triangulation are denoted by boldface letters.

We say a valid triangulation \mathcal{T} is *consistently oriented*, when the geometric realizations of all cells have the same geometric orientation. The orientation of a cell induces a notion of orientation on all of its contained k -sub-simplices for $k = 1, 2$. A k -sub-simplex is called positively oriented if it is positively oriented in the k -dimensional hyperplane bounding the enclosing $(k + 1)$ -sub-simplex with outward pointing normal. This means in particular, that all faces of a cell are positively oriented with respect to the half-plane bounding the cell with a normal pointing to the outside of the cell. If the vertices of a cell are not affinely independent, it is called *degenerate*, and if a cell or any of its k -sub-simplices are not positively oriented, we call it *inconsistently oriented*.

Border faces are faces of a triangulation \mathcal{T} that are incident to only one cell in \mathcal{T} , all other faces are called *inner faces*. Analogously, *border edges* are incident to only one inner face, all other edges are called *inner edges*. The *border of a triangulation* \mathcal{T} is the set of all its border faces. If \mathcal{T} is a valid, consistently oriented triangulation, the geometric realization of its border is a 2-manifold.

The *boundary* ∂S of a subset S of a manifold M are the points in S for which every ε -ball in M contains points in $M \setminus S$. Note that the term *border* is an attribute of the connectivity of a triangulation, whereas *boundary* is a property of its geometric realization.

We need to provide some definitions concerning spherical projections, which we will use to establish geometric conditions for allowable swap operations.

Definition 1. The *spherical projection* of a point $\mathbf{p} \in \mathbb{R}^3$ onto the sphere $S^{\mathfrak{q}}$ with center $\mathbf{q} \in \mathbb{R}^3$ and radius r is defined as

$$\Pi^{\mathfrak{q}}(\mathbf{p}) = \mathbf{q} + r(\mathbf{p} - \mathbf{q})/\|\mathbf{p} - \mathbf{q}\|_2, \quad \mathbf{p} \neq \mathbf{q}.$$

A projection of a set of points $P \subset \mathbb{R}^3 \setminus \{\mathbf{q}\}$ is the set of the projected points,

$$\Pi^{\mathfrak{q}}(P) = \{\Pi^{\mathfrak{q}}(\mathbf{p}) | \mathbf{p} \in P\}.$$

Some properties of the spherical projection (without proof) are:

- If P is a line, $\Pi^{\mathfrak{q}}(P)$ is either two antipodal points (for $\mathbf{q} \in P$), or a half great circle (for $\mathbf{q} \notin P$) of $S^{\mathfrak{q}}$.
- If P is a plane, $\Pi^{\mathfrak{q}}(P)$ is either a great circle (for $\mathbf{q} \in P$), or an open half sphere (for $\mathbf{q} \notin P$) of $S^{\mathfrak{q}}$.
- If $P = \text{conv}(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3)$ is a triangle and the plane defined by P does not contain \mathbf{q} , $\Pi^{\mathfrak{q}}(P)$ is a spherical triangle, bounded by the projection of the edges $\Pi^{\mathfrak{q}}(\text{conv}(\mathbf{p}_1, \mathbf{p}_2))$, $\Pi^{\mathfrak{q}}(\text{conv}(\mathbf{p}_2, \mathbf{p}_3))$, $\Pi^{\mathfrak{q}}(\text{conv}(\mathbf{p}_3, \mathbf{p}_1))$, which are segments of great circles of $S^{\mathfrak{q}}$.

4 Geometric Conditions for \mathcal{G}_{32}

For the geometric conditions to be satisfied for a \mathcal{G}_{32} -swap as defined in Section 2 we have an edge $e = (a, b)$ with n incident cells that is swapped. The triangulation before and after the \mathcal{G}_{32} -swap is denoted by \mathcal{T} and \mathcal{T}' .

Condition 1 *The triangulation $\mathcal{T} = (\mathcal{V}, \mathcal{E}, \mathcal{F}, \mathcal{C})$ is valid, and all cells of \mathcal{T} have positive orientation.*

Condition 2 *The edge e is an inner edge of \mathcal{T} , i.e., every face f incident to e is incident to exactly two cells $c_{f,1} \neq c_{f,2}$.*

Note that the last condition implies that e is not on the border of \mathcal{T} . Furthermore, these conditions induce an order of the faces incident to e .

Lemma 1. *All faces containing e can be ordered to form a cyclic sequence $G = (g_1, \dots, g_n)$, i.e., the index $i = 1, \dots, n$ of g_i is understood modulo n . Furthermore, the dihedral angles θ_i between \mathbf{g}_i and \mathbf{g}_{i+1} (in the direction a to b) are in the interval $(0, \pi)$, and sum to 2π .*

Proof. Due to Condition 2, a face $g = (a, b, v)$ incident to e is incident to two cells $c_{g,1}, c_{g,2}$. Both have two faces incident to e , one of the two is g , the other ones are g'_1 and g'_2 , respectively. The successor of g is the face g'_k of cell $c_{g,k}$ on the positive side of g (in the direction a to b), $k = 1, 2$. The predecessor of g is the other face. Due to Condition 2 this relation determines a cyclic successor-graph without branches.

The dihedral angle θ_i between a face \mathbf{g}_i and its successor \mathbf{g}_{i+1} is the dihedral angle at e of the cell that contains both faces. Therefore, $0 < \theta < \pi$, because otherwise the cell would be inverted or degenerate, contradicting Condition 1.

Since the sequence of faces is cyclic, it surrounds \mathbf{e} . It can only cycle exactly once around \mathbf{e} , because otherwise cells between the faces would intersect in their interior, which contradicts Condition 1. The sum of the dihedral angles between the faces is therefore 2π . \square

We denote the cell between g_i and g_{i+1} as c_i , and the third vertex of g_i as v_i . Thus, Lemma 1 induces also a cyclic order on the cells $C = (c_1, \dots, c_n)$ and vertices $V = (v_1, \dots, v_n)$ around e . Because \mathcal{G}_{32} replaces the cells of C by other cells, we call \mathbf{C} the *affected region*, and the border faces of it are given by

$$\partial\mathbf{C} := \{(a, v_2, v_1), (b, v_1, v_2), \dots, (a, v_1, v_n), (b, v_n, v_1)\},$$

i.e., $\partial\mathbf{C} = \bigcup_{f \in \partial\mathbf{C}} \mathbf{f}$. The line through a and b is denoted by

$$\mathbf{l} = \{\mathbf{a} + \lambda(\mathbf{b} - \mathbf{a}) \mid \lambda \in \mathbb{R}\}. \quad (1)$$

Lemma 2. *There is a closed loop of edges $B = \{b_1, \dots, b_n\}$ that winds around \mathbf{l} exactly once.*

Proof. This follows from Lemma 1, where b_i is the edge of c_i opposite to e . \square

For the sphere $S^{\mathbf{a}}$ around \mathbf{a} contained in the convex hull of all cells containing \mathbf{a} we set $t_1^{\mathbf{a}} = \Pi^{\mathbf{a}}(\mathbf{b})$ and denote by $t_0^{\mathbf{a}}$ the antipodal point of $t_1^{\mathbf{a}}$. Let $\mathbf{B}^{\mathbf{a}} = \Pi^{\mathbf{a}}(\mathbf{B})$ the spherical projection of \mathbf{B} onto $S^{\mathbf{a}}$. Since $\mathbf{B}^{\mathbf{a}}$ is a closed loop on $S^{\mathbf{a}}$, it splits $S^{\mathbf{a}}$ into two parts $S_1^{\mathbf{a}}$ and $S_0^{\mathbf{a}}$, which are characterized by $t_1^{\mathbf{a}} \in S_1^{\mathbf{a}}$ and $t_0^{\mathbf{a}} \in S_0^{\mathbf{a}}$, see Figure 5. Analogously, $S^{\mathbf{b}}$, $t_1^{\mathbf{b}}$, $t_0^{\mathbf{b}}$, $\mathbf{B}^{\mathbf{b}}$, $S_1^{\mathbf{b}}$, and $S_0^{\mathbf{b}}$ are defined.

Definition 2. A *partition* of B is a set $F = \{f_1, \dots, f_m\}$ of faces $f_i \notin \mathcal{F}$, where

1. all vertices of f_i belong to edges of B , i.e., $f_i \subset V$,
2. all edges of f_i are either edges of B or inner edges I , and
3. (a) every edge of B is incident to exactly one face of F ,
(b) every edge of I is incident to exactly two faces of F .

Lemma 3. *Every partition F of B has $n - 3$ inner edges and $m = n - 2$ faces.*

Proof. As a consequence of Lemma 1 partitioning B is equivalent to a triangulation of a simple polygon \mathbf{B}' in a plane perpendicular to \mathbf{l} without introducing new vertices. This polygon is the orthogonal projection of \mathbf{B} along direction \mathbf{l} . Since every simple polygon with n vertices can be triangulated with $n - 2$ triangles (see [2]), i.e., $n - 3$ inner edges, the claim follows. \square

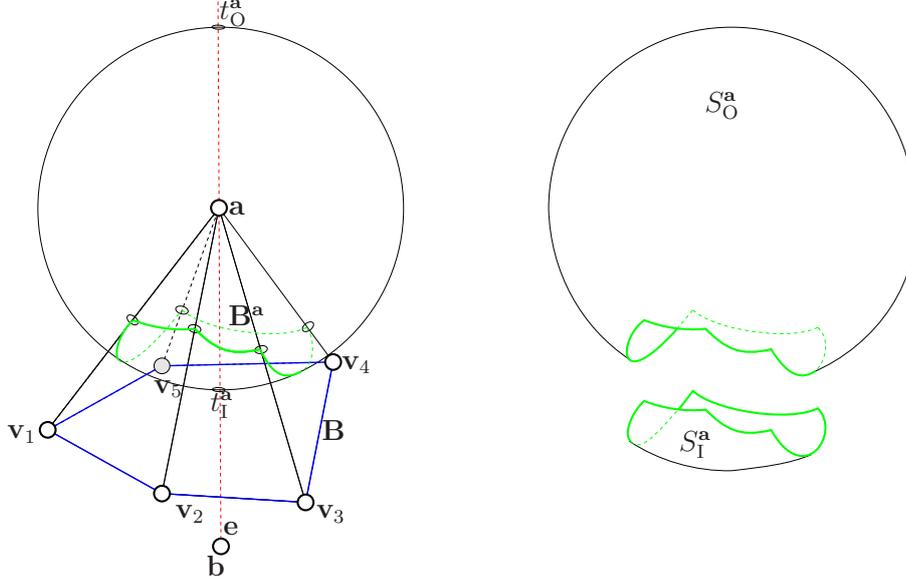


Fig. 5. Terms used in spherical projection with \mathbf{B} in blue and \mathbf{B}^a in green.

The partition F of B defines the cells that are created by the \mathcal{G}_{32} -swap. Every face of the partition is connected to a and b to form two new cells. The set of new cells is $C' = \{c'_{a,1}, c'_{b,1}, \dots, c'_{a,m}, c'_{b,m}\}$ with $c'_{a,j} = f_j \cup \{a\}$ and $c'_{b,j} = \overline{f_j} \cup \{b\}$ for $f_j \in F$. Note that for $n > 3$ the partitions and also the \mathcal{G}_{32} -swap is not unique.

It can happen that C' contains inconsistently oriented or degenerate cells. Therefore, the \mathcal{G}_{32} -swap would result in an invalid triangulation and must not be applied. Whether this happens depends on e and B and also on the choice of F . We call F a *valid partition* if all cells in C' are valid.

Depending on the geometry, there are three different cases. For every case we present an example for $n = 4$, so that two different partitions exist: $F_1 = \{(v_1, v_2, v_3), (v_1, v_3, v_4)\}$ and $F_2 = \{(v_1, v_2, v_4), (v_2, v_3, v_4)\}$. For every example, $\mathbf{a} = (0, 0, 1)$ and $\mathbf{b} = (0, 0, -1)$. Furthermore, the x and y coordinates of \mathbf{v}_1 to \mathbf{v}_4 are $(-0.3, -0.3)$, $(0.7, -1.3)$, $(1.7, -0.3)$, and $(0.7, 0.7)$, respectively.

Every partition is valid For every partition F , all cells in C' are valid. For our example, we choose the z coordinates to be $z_1 = z_2 = z_3 = z_4 = 0$. Both partitions F_1 and F_2 are valid in this case. Note, that every partition is valid as long as the affected region \mathbf{C} is convex, which is only the case if (as in this example) all \mathbf{v}_i are coplanar. But also for a non-convex affected region all partitions can be valid.

Some partitions are invalid For some partitions, there are cells in C' that are inverted or degenerate. But other partitions are valid. For a concrete example, set the z coordinates to $z_1 = z_2 = z_3 = 0.8$ and $z_4 = -0.8$. Here,

F_1 is an invalid partition, because the cell $(\mathbf{a}, \mathbf{v}_1, \mathbf{v}_3, \mathbf{v}_4)$ is inverted, while partition F_2 is valid.

All partitions are invalid It can also happen that no valid partition exists at all. In this case, \mathcal{G}_{32} cannot be applied to e . An example for this case is $z_1 = z_3 = 0.8$ and $z_2 = z_4 = -0.8$. Here, F_1 is invalid because of the inverted cell $(\mathbf{b}, \mathbf{v}_2, \mathbf{v}_4, \mathbf{v}_3)$, F_2 is invalid because of the inverted cell $(\mathbf{a}, \mathbf{v}_1, \mathbf{v}_3, \mathbf{v}_4)$.

These examples show that we need another condition that ensures that F is a valid partition. Under the assumption that Conditions 1 and 2 are satisfied, we found four equivalent formulations 3.2., 3.1., 3.3., and 3.4. for the missing condition. Will prove their equivalence later in Theorem 2. Before we describe the missing condition in detail we need to define the *supporting plane* $\text{pl}(t)$ of a triangle t as the affine hull of its vertices.

Condition 3

- 3.1. All cells $\mathbf{c}'_{a,j}$ and $\mathbf{c}'_{b,j}$ have positive orientation.
- 3.2. Every \mathbf{f}_i has \mathbf{a} on its positive side, and \mathbf{b} on its negative side.
- 3.3. The spherical projection of the faces \mathbf{f}_i onto $S^{\mathbf{a}}$ is contained in $S^{\mathbf{a}} \cup \mathbf{B}^{\mathbf{a}}$, and the interior of the inner edges is projected into $S^{\mathbf{a}}$ (for $S^{\mathbf{b}}$ analogously),

$$\Pi^{\mathbf{P}}(\mathbf{f}_i) \subset S^{\mathbf{P}} \cup \mathbf{B}^{\mathbf{P}}, \quad \text{for all } i = 1, \dots, n, \quad \text{for } \mathbf{P} \in \{\mathbf{a}, \mathbf{b}\}. \quad (2)$$

$$\Pi^{\mathbf{P}}(\mathring{\mathbf{d}}) \subset S^{\mathbf{P}}, \quad \text{for all } d \in I, \quad (3)$$

- 3.4. The interior of the inner edges is a subset of the interior of the affected region, and the supporting planes of all faces f_i intersects the line \mathbf{l} in the interior of \mathbf{e} ,

$$\mathring{\mathbf{d}} \subset \mathbf{C} \setminus \partial \mathbf{C}, \quad \text{for all } d \in I, \quad (4)$$

$$\text{pl}(\mathbf{f}_i) \cap \mathbf{l} \in \mathring{\mathbf{e}}. \quad (5)$$

Theorem 1. If Conditions 1, 2, and 3 are met, the triangulation $\mathcal{T}' = (\mathcal{V}, \mathcal{E}', \mathcal{F}', \mathcal{C}')$ with $\mathcal{C}' = (\mathcal{C} \setminus \mathcal{C}) \cup \mathcal{C}'$ (and \mathcal{E}' and \mathcal{F}' accordingly) is valid.

Proof. Due to Conditions 1 and 3.1., all cells of \mathcal{C}' have positive orientation. To prove that there are no holes in \mathcal{C}' , we check for border faces of the cells of \mathcal{C}' :

- The faces $b_i \cup \{p\}$ for $p \in \{a, b\}$ are border faces of both \mathcal{C} and \mathcal{C}' .
- The faces f_j are incident to $\mathbf{c}'_{a,j}$ and $\mathbf{c}'_{b,j}$, i.e., f_j is not on the border of \mathcal{C}' .
- For the faces $f = d \cup \{p\}$, $d \in I$, $p \in \{a, b\}$, the edge d is incident to two faces f_j and f_k , i.e., f is incident to $\mathbf{c}'_{p,j}$ and $\mathbf{c}'_{p,k}$. So, f is not on the border of \mathcal{C}' .

Thus, there are no new border faces, i.e., there are no holes in \mathcal{C}' . \square

Lemma 4. Condition 3.1. and Condition 3.2. are equivalent.

Proof. By definition, \mathbf{a} is on the positive side of \mathbf{f}_i if and only if the cell $\mathbf{c}'_{a,i}$ has positive orientation. Furthermore, \mathbf{b} is on the negative side of \mathbf{f}_i if and only if the cell $\mathbf{c}'_{b,i}$ has positive orientation. \square

Lemma 5. *Conditions 3.1. and 3.2. imply Condition 3.3.*

Proof. To prove (2) we first show that $\Pi^{\mathbf{a}}(\mathbf{F})$ is a connected region on $S^{\mathbf{a}}$ that is bounded by $\mathbf{B}^{\mathbf{a}}$. Then we show that $t_1^{\mathbf{a}} \in \Pi^{\mathbf{a}}(\mathbf{F})$.

Due to Condition 3.1. \mathbf{a} is not in \mathbf{F} , since this would cause degenerate cells, and $\Pi^{\mathbf{a}}(\mathbf{F})$ is a connected region on $S^{\mathbf{a}}$. For $f_i, f_j \in F$ with common edge $d \in I$, the spherical triangles $\Pi^{\mathbf{a}}(\mathbf{f}_i)$ and $\Pi^{\mathbf{a}}(\mathbf{f}_j)$ share the spherical edge $\Pi^{\mathbf{a}}(\mathbf{d})$. Due to Condition 3.1. the both cells $\mathbf{c}'_{a,i}$ and $\mathbf{c}'_{a,j}$ have positive orientation, so they are on opposite sides of the plane P through \mathbf{d} and \mathbf{a} . Therefore, $\Pi^{\mathbf{a}}(\mathbf{f}_i)$ and $\Pi^{\mathbf{a}}(\mathbf{f}_j)$ are also on opposite sides of $\Pi^{\mathbf{a}}(\mathbf{d})$, see Figure 6. This implies that the interior of all inner edges of I is not projected to the boundary of $\Pi^{\mathbf{a}}(\mathbf{F})$. The same holds true for all interior points of \mathbf{F} . Thus, the boundary of $\Pi^{\mathbf{a}}(\mathbf{F})$ consists of projections of the border edges of B . Consequently, the interior of $\Pi^{\mathbf{a}}(\mathbf{F})$ is not intersected by $\mathbf{B}^{\mathbf{a}}$, so $\Pi^{\mathbf{a}}(\mathbf{F})$ is either completely in $S_1^{\mathbf{a}} \cup \mathbf{B}^{\mathbf{a}}$, or in $S_0^{\mathbf{a}} \cup \mathbf{B}^{\mathbf{a}}$.

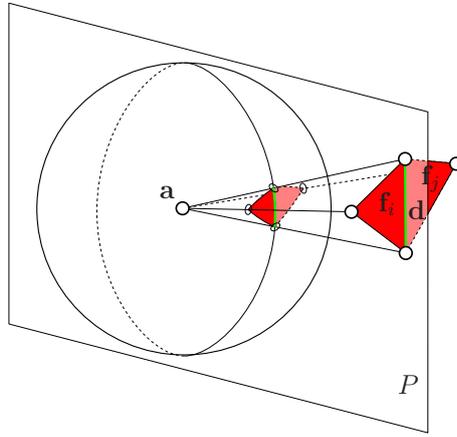


Fig. 6. The projections of \mathbf{f}_i and \mathbf{f}_j are on opposite sides of the projection of the edge \mathbf{d} .

Since \mathbf{B} winds around \mathbf{l} once, the \mathbf{l} line intersects \mathbf{F} in at least one face \mathbf{f}_i . Let $\mathbf{p} = \mathbf{l} \cap \mathbf{f}_i$. Because \mathbf{a} is on the positive side of \mathbf{f}_i (Condition 3.2.), $\Pi^{\mathbf{a}}(\mathbf{p}) = t_1^{\mathbf{a}}$. Therefore, $\Pi^{\mathbf{a}}(\mathbf{F}) \subset S_1^{\mathbf{a}} \cup \mathbf{B}^{\mathbf{a}}$. Analogously, one can show $\Pi^{\mathbf{b}}(\mathbf{F}) \subset S_1^{\mathbf{b}} \cup \mathbf{B}^{\mathbf{b}}$.

Especially, the interior of $\Pi^{\mathbf{a}}(\mathbf{d})$ does not intersect $\mathbf{B}^{\mathbf{a}}$, which implies (3). \square

Lemma 6. *Condition 3.3. implies Condition 3.4.*

Proof. Let $d \in I$ be an inner edge of F , and $\mathbf{p} \in \mathring{\mathbf{d}}$ be an interior point of \mathbf{d} . Due to Condition 3.3., $\mathbf{p}^{\mathbf{a}} = \Pi^{\mathbf{a}}(\mathbf{p}) \in S_1^{\mathbf{a}}$. We split $S_1^{\mathbf{a}}$ into spherical triangles by adding edges from $\Pi^{\mathbf{a}}(\mathbf{v}_i)$ to $t_1^{\mathbf{a}}$. At least one of these triangles contains $\mathbf{p}^{\mathbf{a}}$. Let this triangle be $t = (t_1^{\mathbf{a}}, \Pi^{\mathbf{a}}(\mathbf{v}_l), \Pi^{\mathbf{a}}(\mathbf{v}_{l+1}))$, see Figure 7. The boundary $\mathbf{B}^{\mathbf{a}}$ (green) is partitioned into spherical triangles (red lines), \mathbf{d} is $(\mathbf{v}_1, \mathbf{v}_4)$

(blue line) and $\mathbf{p} \in \mathring{\mathbf{d}}$. In this case, $\Pi^{\mathbf{a}}(\mathbf{p})$ is within the spherical triangle $t = (t_1^{\mathbf{a}}, \Pi^{\mathbf{a}}(\mathbf{v}_4), \Pi^{\mathbf{a}}(\mathbf{v}_5))$.

The set of points that are projected into t is defined as the intersection of the half spaces defined by the planes spanned by \mathbf{a} and one of the edges of t , i.e., $g_1 = (a, b, v_l)$, $g_2 = (a, v_{l+1}, b)$, and $g_3 = (a, v_l, v_{l+1})$, which contains the fourth point $\{a, b, v_l, v_{l+1}\} \setminus g_i$. The point \mathbf{p} cannot be on the negative side of \mathbf{g}_1 , \mathbf{g}_2 or \mathbf{g}_3 , as this would mean that its image is not in t . Also, it cannot be in the plane defined by \mathbf{a} and \mathbf{g}_3 , as this would mean that it is projected to $\mathbf{B}^{\mathbf{a}}$.

With the same argument for $\Pi^{\mathbf{b}}$, we obtain the faces $g_4 = (b, v_l, a)$, $g_5 = (b, v_{l+1}, a)$, and $g_6 = (b, v_{l+1}, v_l)$. Removing the redundant faces $g_4 \equiv g_1$ and $g_5 \equiv g_2$, we can conclude that \mathbf{p} is not on the negative side of \mathbf{g}_1 and \mathbf{g}_2 , and it is on the positive side of \mathbf{g}_3 and \mathbf{g}_6 . These four faces define the cell c_i . Thus, $\mathbf{p} \in \mathbf{C}$, and $\mathbf{p} \notin \partial\mathbf{C}$, proving (4).

We still have to prove (5). Assume there exists a face f in F with $\{\mathbf{q}\} = \text{pl}(f) \cap \mathbf{l} \notin \mathring{\mathbf{e}}$ and, without loss of generality, $\lambda \leq 0$. This face has at least one interior edge $d \in I$ and we chose an arbitrary point $\mathbf{p} \in \mathring{\mathbf{d}}$. Now, \mathbf{p} is projected to $\mathbf{p}^{\mathbf{a}}$ which lies outside of $S_1^{\mathbf{a}}$. This contradicts (3) and, thus, proves (5). \square

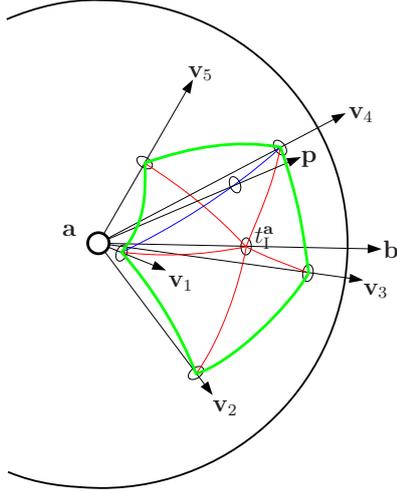


Fig. 7. $S_1^{\mathbf{a}}$ is divided into spherical triangles (red lines), one of which contains $\Pi^{\mathbf{a}}(\mathbf{p})$.

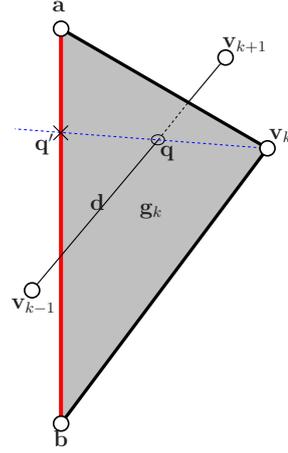


Fig. 8. The intersection of the extension of segment \mathbf{v}_k to \mathbf{q} with \mathbf{l} is between \mathbf{a} and \mathbf{b} .

Lemma 7. *If Conditions 1 and 2 are satisfied, Condition 3.4. implies Condition 3.2.*

Proof. For $n = 3$ we have $I = \emptyset$ and $F = \{f_1\}$. Since B circles around \mathbf{l} , there must be an intersection of \mathbf{l} and \mathbf{f}_1 . Due to Condition 3.4., this is between \mathbf{a} and

\mathbf{b} , and because of the order of the vertices of f_1 as induced by Lemma 1, \mathbf{a} is on the positive and \mathbf{b} on the negative side of \mathbf{f}_1 , and Condition 3.2. is satisfied.

We now consider $n > 3$. The partition F contains $n - 2$ faces, the border B has n edges (see Lemma 3). If every face of F had at most one edge of B , there would be at least two edges in B left. Since no face of F can have three edges of B (otherwise B would have a sub-cycle of three edges), at least two faces of F must have two edges of B . Let $\hat{F} \subset F$ be the set of faces with two edges in B .

The line \mathbf{l} intersects either one face of F in its interior, or it intersects an inner edge of I and therefore two faces of F on their border.

In the case that \mathbf{l} intersects an inner edge, and the adjacent faces of F are the only two faces in \hat{F} , there can be no other faces in F , due to the following: if two faces with each two edges in B and both sharing a common inner edge, their edges in B already define a cycle. Since B does not contain any sub-cycles, there can be no further edges in B . In this case we have $n = 4$. Since the intersection of \mathbf{l} with \mathbf{f}_1 and \mathbf{f}_2 is between \mathbf{a} and \mathbf{b} (Condition 3.4.), and because of the order of the vertices of f_1 and f_2 , \mathbf{a} is on the positive side of \mathbf{f}_1 and \mathbf{f}_2 , and \mathbf{b} is on the negative side. Thus, in this case Condition 3.2. is satisfied.

For the remaining case there is at least one face f in \hat{F} that has no intersection with \mathbf{l} , because otherwise Conditions 1 and 2 were violated. Let $f = (v_{k-1}, v_k, v_{k+1})$. Because f does not intersect \mathbf{l} , $\theta_{k-1} + \theta_k < \pi$. Thus, the inner edge $d = (v_{k-1}, v_{k+1})$ cannot cross any other cell besides c_{k-1} and c_k . Due to Condition 3.4., $\mathbf{d} \subset c_{k-1} \cup c_k$. With $g_k = (a, b, v_k)$, the intersection $\mathbf{d} \cap \mathbf{g}_k = \{\mathbf{q}\}$, with \mathbf{q} in $\mathring{\mathbf{g}}_k$. When extending the line segment from \mathbf{v}_k to \mathbf{q} , it intersects the segment \mathbf{e} in its interior in point \mathbf{q}' , because of (5) (see Figure 8). Since \mathbf{v}_k and \mathbf{q} are points in \mathbf{f} , the line through \mathbf{v}_k and \mathbf{q} is also in the plane of \mathbf{f} , and so is \mathbf{q}' . From these considerations and the vertex order of f , it follows that \mathbf{a} is on the positive and \mathbf{b} is on the negative side of \mathbf{f} . Thus, \mathbf{f} fulfills Condition 3.2..

Now we remove f from F , i.e., F becomes $F \setminus \{f\}$, B becomes $(B \setminus \{b_{k-1}, b_k\}) \cup \{d\}$, and I becomes $I \setminus \{d\}$. This new edge cycle B still satisfies Conditions 1 and 2, but has one edge less. This procedure can be repeated until $n = 3$, or $n = 4$ and \mathbf{l} intersects both faces in F . \square

Theorem 2. *If Conditions 1 and 2 are satisfied, Conditions 3.2., 3.1., 3.3., and 3.4. are equivalent.*

Proof. This follows directly from Lemmata 4, 5, 6, and 7. \square

5 Geometric Conditions for \mathcal{G}_{23}

We use the same notation as in Section 4, i.e., $F = \{f_1, \dots, f_m\}$ is a set of faces sandwiched between \mathbf{a} and \mathbf{b} , such that \mathbf{F} is a connected 2-manifold. The edge set $B = \{b_1, \dots, b_n\}$ are the border edges of F . The order of edges in B is induced by the order of boundary edges in $\partial\mathbf{F}$.

The triangulation before and after the \mathcal{G}_{23} -swap is denoted by T' and T . We define the orientation of \mathbf{f}_i so that \mathbf{a} is on the positive side of \mathbf{f}_i . The cells

incident to these faces are $C' = \{c'_{a,1}, c'_{b,1}, \dots, c'_{a,m}, c'_{b,m}\}$ with $c'_{p,i} = f_i \cup \{p\}$ for $i = 1, \dots, m$ and $p \in \{a, b\}$. The new edge in \mathcal{T} is $e = (a, b)$.

Next we define the conditions for which \mathcal{G}_{23} will result in a valid triangulation.

Condition 4 *The triangulation $\mathcal{T}' = (\mathcal{V}, \mathcal{E}', \mathcal{F}', C')$ is valid, and all cells of \mathcal{T}' have positive orientation.*

Condition 5 *The edges of B form exactly one simple cycle (v_1, \dots, v_n) .*

This condition ensures that the faces in F are connected via edges, that there is only one connected component of faces, and that the faces form a bounded 2-manifold without holes. Examples of sets F that violating Condition 5 are shown in Figures 9(a), 9(b), and 9(c).

Condition 6 *All vertices incident to a face in F are on the border B .*

Condition 6 the absence of interior vertices in F that are not part of B . Those interior vertices would be removed by \mathcal{G}_{23} , but a swap may only modify the connectivity, but not add, remove, or move vertices. Figure 9(d) shows an example of a set F that violates Condition 6 due to an interior vertex.

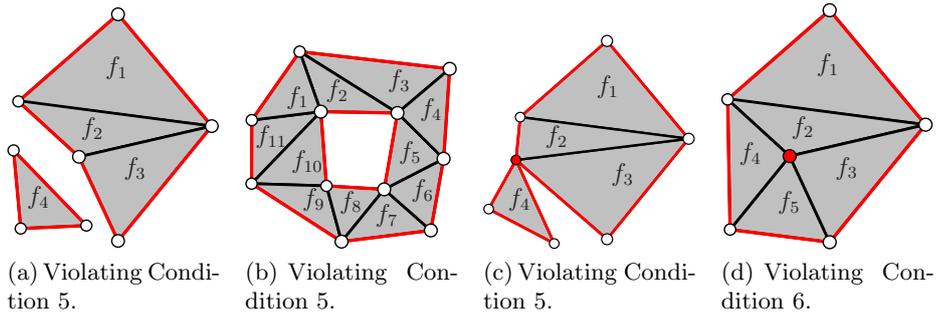


Fig. 9. Examples of sets F violating Conditions 5 or 6.

Lemma 8. *If Condition 6 is satisfied, the number of vertices in B is $n = m + 2$.*

Proof. If Condition 6 is satisfied, F is a partition of B . Considering Lemma 3 we can conclude $m = n - 2$. Therefore, $n = m + 2$. \square

The \mathcal{G}_{23} -swap will now replace the cells C' by the cells $C = \{c_1, \dots, c_n\}$ with

$$c_i = (a, b, v_i, v_{i+1})$$

and faces $g_i = (a, b, v_i)$, where the index i is understood modulo n .

Condition 7 *One of the equivalent following conditions holds:*

7.1. All cells \mathbf{c}_i have positive orientation.

7.2. The dihedral angle θ_i between the faces \mathbf{g}_i and \mathbf{g}_{i+1} (in counterclockwise direction, seen from a in direction b) is in $(0, \pi)$.

Lemma 9. Condition 7.1. and Condition 7.2. are equivalent.

Proof. \mathbf{c}_i has positive orientation if and only if \mathbf{c}_i is consistently oriented or non-degenerate. This is true if and only if the inner dihedral angle θ_i is in $(0, \pi)$. \square

Theorem 3. If Conditions 4, 5, 6, and 7 are met, the triangulation $\mathcal{T} = (\mathcal{V}, \mathcal{E}, \mathcal{F}, \mathcal{C})$ with $\mathcal{C} = (\mathcal{C} \setminus \mathcal{C}') \cup \mathcal{C}$ (and \mathcal{E} and \mathcal{F} accordingly) is valid.

Proof. Due to Conditions 4 and 7.1., all cells of \mathcal{C} have positive orientation. To prove that there are no holes in \mathcal{C} , we check for border faces of the cells of \mathcal{C} :

- The faces $b_i \cup \{p\}$ for $p \in \{a, b\}$ are border faces of both \mathcal{C}' and \mathcal{C} .
- The faces g_i are incident to c_{i-1} and c_i , i.e., g_i is not on the border of \mathcal{C} .

Thus, there are no new border faces, i.e., there are no holes in \mathcal{C} . \square

6 Replacing Generalized Swaps by a Series of Basic Swaps

In [23] Shewchuk showed that the “multi-face removal” (equivalent to \mathcal{G}_{23}) and “edge removal” (equivalent to \mathcal{G}_{32}) can be replaced by a series of basic 2-3 and 3-2 swaps. The intermediate triangulations are topologically correct, but may contain inconsistently oriented or degenerate tetrahedra.

We will show that there always exists a series of basic 2-3, 3-2, and 4-4 swaps to mimic the effect of a \mathcal{G}_{23} - and a \mathcal{G}_{32} -swap, where all intermediate triangulations are valid. This result shows that the \mathcal{G}_{23} - and \mathcal{G}_{32} -swaps do not add additional potential that is not already possible with 2-3, 3-2 and 4-4 swaps. An optimization procedure like simulated annealing should theoretically be able to find a near-optimal solution also without utilizing \mathcal{G}_{23} and \mathcal{G}_{32} . In practice, the convergence rate can be increased by implementing \mathcal{G}_{23} and \mathcal{G}_{32} .

6.1 Replacing \mathcal{G}_{32}

Let \mathbf{e} be an inner edge of triangulation T , B the set of border edges, and F a valid partition of B , so that the Conditions 1, 2, and 3 for \mathcal{G}_{32} are satisfied.

Theorem 4. The same effect as the \mathcal{G}_{32} swap operation of \mathbf{e} and partition F can be obtained by a series of either

- $n - 3$ basic 2-3 swaps followed by a 3-2 swap, or
- $n - 4$ basic 2-3 swaps followed by a 4-4 swap, for $n \geq 4$.

Proof. We use the same arguments as in the proof of Lemma 7.

For $n = 3$, F consists of exactly one face f_1 , and the vertices of f_1 circle around \mathbf{e} exactly once. Therefore, the conditions are satisfied to apply a 3-2 swap to \mathbf{e} , so we can substitute \mathcal{G}_{32} by a single 3-2 swap.

For $n = 4$, and the single inner edge $d = (v_i, v_{i+2})$ with $i \in \{1, 2\}$ intersects with \mathbf{e} , the quadrilateral $(\mathbf{v}_i, \mathbf{a}, \mathbf{v}_{i+2}, \mathbf{b})$ is planar and convex, fulfilling the conditions of a 4-4 swap. This 4-4 swap replaces e by d and the four cells of C with the four cells of C' . Thus, the \mathcal{G}_{32} swap can be replaced by a single 4-4 swap.

If $n = 4$ and d and e do not intersect, or if $n > 4$, there is at least one face in F with two edges in B that does not intersect \mathbf{e} . Let this face be $f_j = (v_{i-1}, v_i, v_{i+1})$. As in the proof for Lemma 7, the edge $d = (v_{i-1}, v_{i+1})$ intersects the face g_i in its interior, so the cells c_{i-1} and c_i fulfill the condition for a 2-3 swap. This swap removes c_{i-1} and c_i from the triangulation, adds $c'_{a,j}$ and $c'_{b,j}$ and a temporary new cell $c = (a, b, v_{i-1}, v_{i+1})$. The remaining cells $(C \setminus \{c_{i-1}, c_i\}) \cup \{c\}$ together with the reduced partition $F \setminus \{f_j\}$ and the reduced border $(B \setminus \{(v_{i-1}, v_i), (v_i, v_{i+1})\}) \cup \{d\}$ fulfill Conditions 1-3. So, \mathcal{G}_{32} can be applied to the reduced setting. By induction, the reduced setting can be processed with either $(n - 1) - 3$ 2-3 swaps, followed by a 3-2 swap, or with $(n - 1) - 4$ 2-3 swaps, followed by a 4-4 swap. Adding the 2-3 swap to remove f_j , the claim follows. \square

6.2 Replacing \mathcal{G}_{23}

Since the \mathcal{G}_{23} operation is the inverse of the \mathcal{G}_{32} operation for the same partition F , \mathcal{G}_{23} can be replaced by a series of basic swaps.

Theorem 5. *The same effect as a \mathcal{G}_{23} operation of a partition F sandwiched between a and b can be obtained by a series of either*

- a single 2-3 swap, followed by $n - 3$ 3-2 swaps, or
- a single 4-4 swap, followed by $n - 4$ 3-2 swaps.

Proof. While \mathcal{G}_{23} replaces the cells C' by cells C , \mathcal{G}_{32} does the inverse. \mathcal{G}_{32} can be substituted by a series of basic swap operations $\mathcal{G}_{32} = s_1 \circ s_2 \circ \dots \circ s_m$, with m being either $n - 2$ (s_m being a 3-2 swap) or $n - 3$ (s_m being a 4-4 swap), as in Theorem 4. For the same choice of F , we have

$$\mathcal{G}_{23} = \mathcal{G}_{32}^{-1} = (s_1 \circ \dots \circ s_m)^{-1} = s_m^{-1} \circ \dots \circ s_1^{-1}.$$

The inverse of a 3-2 swap is a 2-3 swap and vice versa, and the inverse of a 4-4 swap is a corresponding 4-4 swap. We start in \mathcal{G}_{23} with s_m^{-1} , which is either a 2-3 swap or a 4-4 swap. Then we proceed with either $n - 3$ or $n - 4$ 3-2 swaps. \square

7 Conclusions

We have presented different geometric conditions for generalized swap operations a 3D triangulation. These conditions are proved to be equivalent, such that one

can use that particular condition in practice that is most appropriate given the specific needs of an implementation. In a mesh optimization application these swap operations are used to speed up the optimization process and to attenuate "getting stuck" in local minima.

Furthermore, we have shown that the generalized swap operations can be realized by simple 3-2, 2-3, and 4-4 swaps, which simplifies the implementation significantly. This decomposition of the generalized swap guarantees at the same time, that all intermediate triangulations are consistently oriented and do not contain degenerate cells, causing numerical problems in certain applications.

Based on these conditions, our future research plans are focused on applications of 3D mesh optimizations, e.g., in video compressions or bio-medical and bio-mechanical simulations.

Acknowledgments

This work was partly funded by the German Research Foundation (DFG) within the International Research Training Group 1131, "Visualization of Large and Unstructured Data Sets" at the University of Kaiserslautern. This work was also supported by the National Science Foundation under contract ACI 9624034 (CAREER Award) and a large Information Technology Research (ITR) grant.

References

1. N. Amenta, S. Choi, and R. K. Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry*, 19(2–3):127–153, 2001.
2. M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry*. Springer, 2nd edition, 2000.
3. H. L. de Cougny and M. S. Shephard. Parallel refinement and coarsening of tetrahedral meshes. *Int. J. Numer. Meth. Engng.*, 46:1101–1125, 1999.
4. L. Demaret, N. Dyn, and A. Iske. Image compression by linear splines over adaptive triangulations. *Signal Process.*, 86(7):1604–1616, 2006.
5. L. Demaret, N. Dyn, A. Iske, and M. Floater. Adaptive thinning for terrain modelling and image compression. In N. A. Dodgson, M. S. Floater, and M. A. Sabin, editors, *Advances in Multiresolution for Geometric Modelling*, pages 321–340. Springer, 2004.
6. Klaus Denker, Burkhard Lehner, and Georg Umlauf. Online triangulation of laser-scan data. In R. Garimella, editor, *Proceedings of the 17th International Meshing Roundtable 2008*, pages 415–432, 2008.
7. N. Dyn, D. Levin, and S. Rippa. Data dependent triangulations for piecewise linear interpolations. *IMA Journal of Numerical Analysis*, 10(1):137–154, Jan 1990.
8. M. Garland and P. Heckbert. Fast polygonal approximation of terrains and height fields. Technical report, CS Department, Carnegie Mellon University, 1995.
9. M. Garland and P. S. Heckbert. Surface simplification using quadric error metrics. In *SIGGRAPH '97*, pages 209–216, 1997.
10. M. Garland and Y. Zhou. Quadric-based simplification in any dimension. *ACM Trans. Graph.*, 24(2):209–239, 2005.
11. H. Hoppe. Progressive meshes. In *SIGGRAPH '96*, pages 99–108, 1996.

12. H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Computer Graphics*, 26(2):71–78, 1992.
13. B. Joe. Construction of three-dimensional improved-quality triangulations using local transformations. *SIAM J. Sci. Comp.*, 16(6):1292–1307, 1995.
14. B. M. Klingner and J. R. Shewchuk. Aggressive tetrahedral mesh improvement. In *16th Int. Meshing Roundtable*, pages 3–23, 2007.
15. O. Kreylos and B. Hamann. On simulated annealing and the construction of linear spline approximations for scattered data. *IEEE TVCG*, 7(1):17–31, 2001.
16. C. L. Lawson. Properties of n -dimensional triangulations. *CAGD*, 3:231–246, 1986.
17. B. Lehner. *Meshing Techniques for Image/Video Compression and Surface Reconstruction*. PhD thesis, TU Kaiserslautern, Germany, 2008.
18. B. Lehner, G. Umlauf, and B. Hamann. Image compression using data-dependent triangulations. In G. Bebis, editor, *ISVC 2007*, pages 351–362, 2007.
19. B. Lehner, G. Umlauf, and B. Hamann. Video compression using data-dependent triangulations. In *Computer Graphics and Visualization '08*, pages 244–248, 2008.
20. H. Pedrini. An improved refinement and decimation method for adaptive terrain surface approximation. In *Proceedings of WSCG*, pages 103–109, 2001.
21. V. Petrovic and F. Küster. Optimized construction of linear approximations to image data. In *Proc. 11th Pacific Conf. on Comp. Graphics and Appl.*, pages 487–491, 2003.
22. L. L. Schumaker. Computing optimal triangulations using simulated annealing. *CAGD*, 10(3-4):329–345, 1993.
23. J. R. Shewchuk. Two discrete optimization algorithms for the topological improvement of tetrahedral meshes, 2002. Unpublished manuscript, <http://www.cs.berkeley.edu/~jrs/papers/edge.pdf>.
24. J. R. Shewchuk. What is a good linear element? Interpolation, conditioning, and quality measures. In *11th Int. Meshing Roundtable*, pages 115–126, 2002.
25. X. Yu, B. S. Morse, and T. W. Sederberg. Image reconstruction using data-dependent triangulation. *IEEE Comput. Graph. Appl.*, 21(3):62–68, 2001.