SERIE B — INFORMATIK

Surface Reconstruction between Simple Polygons via Angle Criteria[◇]

Emo Welzl*
Barbara Wolfers*

B 94-11
April 1994

Abstract

We consider the problem of connecting two simple polygons P and Q in parallel planes by a polyhedral surface. The goal is to find an optimality criterion which naturally satisfies the following conditions: (i) if P and Q are convex, then the optimal surface is the convex hull of P and Q (without facets P and Q), and (ii) if P can be obtained from Q by scaling with a center c, then the optimal surface is the portion of the cone defined by P and apex c between the two planes. We provide a criterion (based on the sequences of angles of the edges of P and Q), which satisfies these conditions, and for which the optimal surface can be efficiently computed. Moreover, we supply a condition, so-called *angle consistency*, which proved very helpful in preventing self intersections (for our and other criteria). The methods have been implemented and gave improved results in a number of examples.

◇Work partially supported
*Institut für Informatik, Fachbereich Mathematik, Freie Universität Berlin,
Arnimallee 2–6, W1000 Berlin 33, Germany. E-mail: welzl@tcs.fu-berlin.de, wolfers@tcs.fu-berlin.de.

1 Introduction

The reconstruction of a three-dimensional object from its cross-sections data is a problem with many applications like clinical medicine (computerized tomography and magnetic resonance imaging), biomedical research, computer graphics, animation, geology, etc., [Sch].

Here is the set-up we want to consider: P and Q are simple polygons in parallel planes h_P and h_Q, respectively. A *surface between P and Q* is a cyclic sequence of triangles, each triangle is the convex hull of an edge of one of the polygons and a vertex of the other polygon; consecutive triangles share an edge (connecting a vertex from P with a vertex from Q), and the sequence encounters the edges of P in the same counterclockwise order as P, and analogously for Q. So we ignore the problems arising from the fact that the cross sections of an object may contain several polygons (polygons have to be assigned to each other, and 'branchings' may occur). This can be handled by a preprocessing step by other methods, see e.g. [MK], [MSS]. Moreover, we restrict ourselves by not allowing other vertices in the surface but those in P and Q.

A number of methods have been proposed in the literature. For example there is the volume based approach [BGLS], [LC], the paper by Barequet and Sharir [BS], and the work by Boissonnat [B], [BG], based on Delaunay triangulation. Most methods associate with every potential connecting surface a parameter (usually a real number), and the surface of choice is one which optimizes (minimizes, maximizes) this parameter. Examples are: (1) surface of minimum area [FKU], [SP], (2) surface where the resulting enclosed solid has maximal volume [K], (3) surface, where the overall edge length is minimal, etc. [WA], [SG]. Other approaches [C], [ChrS], [GD] start the construction at some point and proceed according to local criteria.

It turns out that these methods have drawbacks, which occur already in simple natural examples: probably most striking is the case of two regular n-gons P and Q, where the orthogonal projection of P in h_Q is sufficiently far apart from Q (the optimal surface according to the minimum area criterion is depicted in Figure 1).

Our starting point was to set up general requirements which should be met by a 'good' optimality criterion in a natural way:

> **Condition C1.** If P and Q are convex polygons, then the optimal surface is the convex hull of P and Q (without facets P and Q).
>
> **Condition C2.** If P can be obtained from Q by scaling with a center c, then the optimal solution is the portion of a cone defined by P with apex c between the two planes h_P and h_Q. Similarly, if P is a translate of Q, then the surface should be a cylindric section.

Surprisingly enough, none of the criteria we found in the literature satisfy both conditions (Figure 1 demonstrates that the minimum area criterion violates both conditions). Our method starts with the following simple observation. The sequence of triangles from a surface defines a 'merge' of the edges from P and Q, (go through the sequence of triangles and for each one take the edge which is from P or Q, see

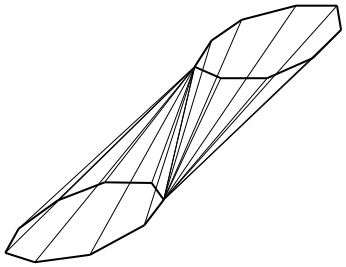

Figure 1: Area-optimal solution for two regular 9-gons.

Figure 2). This sequence yields again a polygon (not necessarily simple!), which has

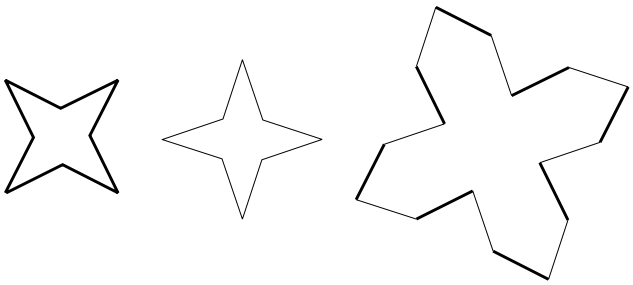

Figure 2: Merge of two polygons.

also a geometric interpretation in terms of the surface: If all the edges are halved in length, then we get the polygon obtained by intersecting the surface with the plane half way between h_P and h_Q. For every such merged polygon we add up the absolute values of the 'turning angles' $\delta(e, e')$ between any pair of consecutive edges e and e'. A surface is called *optimal* if its associated polygon-merge minimizes this sum. The intuition is that we try to keep the surface (or, more precisely, its intersection with planes parallel to h_P) as smooth as possible.

In this way we satisfy conditions C1 and C2, as we will prove in Section 2. It may appear to be more appropriate to consider the sum of squares of $\delta(e, e')$ instead, but, as it turns out, this violates condition C2.

There is the issue of surfaces with self-intersections – definitely an undesired effect – which we have not touched so far. This may very well happen for the optimal surfaces (also for our criterion). As a matter of fact, Gitlin, O'Rourke and Subramanian [GORS], show that there are instances of polygons which do not allow a connecting surface (in the way we defined it) without self-intersections (one polygon may even be chosen as a triangle). (There is a subtle issue what we call a self-intersection, but we do not elaborate on this; e.g. the surface in Figure 1 has a self-intersection in the sense of [GORS].)

Section 3 describes the so-called *angle-consistency condition* for merged polygons. Roughly speaking, this disallows that in the merged sequence between two edges in P there is a sequence of edges in Q which runs into a spiral without 'resolving' it. Experiments show, that the condition prevents self-intersections in many examples, and we prove that a violation of the condition *enforces* a self-intersection (i.e. requiring angle-consistency does not exclude any good solutions).

The algorithmic aspects are dealt with in Section 4. We show that the optimal angle-consistent solution with respect to our angle criterion can be computed in time $O((dt)^4 + m + n)$, where m and n are the numbers of edges of P and Q, d is a parameter that indicates to what extent P or Q run into spirals, and t counts the number of edges of inflection in P and Q, (i.e. edges where preceding and succeeding vertex lie on opposite sides of the line through the edge; e.g., for a convex polygon this parameter is 0). In many instances, d and t are very small compared to the number of edges.

We have implemented our method, and some other methods for the sake of comparison. The angle-consistency condition has been directly motivated by phenomena we observed on results of the implementation in simple natural examples.

Clearly, the 'best' surface will always depend on the specific application, and there may even occur applications where our conditions C1 and C2 are not appropriate. Nevertheless, we believe that our method represents an interesting alternative to the existing ones. Moreover, merged polygons raise some mathematically interesting questions. We refer to [GRS] for a paper treating some related aspects.

2 An angle criterion for merging polygons.

We first introduce some simple notation for sequences and polygons.

Notation for sequences. Given two sequences $X = (x_0, x_1, \ldots, x_{n-1})$ and $Y = (y_0, y_1, \ldots, y_{m-1})$, we say that X and Y are *cyclically equivalent*, denoted by $X =_{\text{cyc}} Y$, if $n = m$ and there exists an i, $0 \le i \le n-1$, such that $(x_i, x_{i+1}, \ldots, x_{n-1}, x_0, x_1, \ldots, x_{i-1}) = (y_0, y_1, \ldots, y_{m-1})$. We adopt the convention that indices are taken modulo the length of the considered sequence, in particular $x_n = x_0$.

Let $Z = (z_0, z_1, \ldots, z_{n-1})$ be a sequence, and let $I = \{i_1, i_2, \ldots, i_k\}$, $0 \le i_1 < i_2 < \cdots i_k \le n-1$. The *$I$-restriction, $Z_{|I}$, of Z* is the sequence $(z_{i_1}, z_{i_2}, \ldots, z_{i_k})$.

Z is a *cyclic merge of* sequences X and Y if there is a partition (I, J) of $\{0, 1, \ldots, n-1\}$ such that $X =_{\text{cyc}} Z_{|I}$ and $Y =_{\text{cyc}} Z_{|J}$. Note that I and J are not uniquely determined; in order to be more specific about which elements come from which sequence, we call Z the *(I, J)-indexed cyclic merge of* X and Y.

Polygons. A *polygon* P is a sequence $(p_0, p_1, \ldots, p_{n-1})$ of $n \ge 2$ points in the plane, such that $p_i \ne p_{i+1}$ for all i, $i = 0, 1, \ldots, n-1$. Two polygons are considered *equivalent* if their defining sequences are cyclically equivalent.

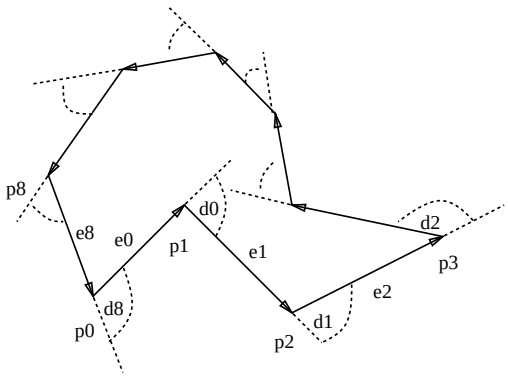

Figure 3: Edge vectors and turning angles, $\delta_0 < 0, \delta_1 > 0$ etc.

A polygon is *simple* if $n \geq 3$, all points p_i, $i = 0, 1, \ldots, n-1$, are pairwise distinct, and each open line segment $\overline{p_i p_{i+1}}$, $i = 0, 1, \ldots, n-1$, is disjoint from all p_j, $j = 0, 1, \ldots, n-1$, and from all $\overline{p_j p_{j+1}}$, $j = 0, 1, \ldots, n-1$, $j \neq i$.

Every polygon $P = (p_0, p_1, \ldots, p_{n-1})$ defines

- a sequence of *edge vectors* $E_P = (e_0, e_1, \ldots, e_{n-1}) \in (\mathbb{R}^2 - o)^n$ where[1] $e_i = p_{i+1} - p_i$, $o = (0,0)$ is the zero vector;
- a sequence of *edge angles* $A_P = (a_0, a_1, \ldots, a_{n-1}) \in (S^1)^n$ where $a_i = e_i / \| e_i \|$;
- and a sequence of *turning angles* $\Delta_P = (\delta(e_0, e_1), \delta(e_1, e_2), \ldots, \delta(e_{n-1}, e_n = e_0)) \in ((-\pi, +\pi) \cup \{\perp\})^n$ where $\delta : ((\mathbb{R}^2 - o) \times (\mathbb{R}^2 - o))^n \longrightarrow ((-\pi, +\pi) \cup \{\perp\})$ and $\delta(e, e')$, is the counterclockwise angle between e and e' in the interval $(-\pi, +\pi)$.

$\delta(e_i, e_{i+1})$, for short δ_i, can be seen as the turn of the tangent at point p_{i+1}, where a counterclockwise turn gives a positive value, and a clockwise turn gives a negative value; see Figure 3. If $a_i = -a_{i+1}$, i.e. e_i and e_{i+1} are oppositely directed, then we define $\delta(e_i, e_{i+1}) = \perp$; intuitively, $\perp$ represents $\pm\pi$. If $\delta(e_i, e_{i+1}) = \perp$ then $|\delta(e_i, e_{i+1})| := \pi$.

Given the edge vector sequence $E_P = (e_0, e_1, \ldots, e_{n-1})$ of a polygon P we write $\delta(P) = \delta(E_P) := \sum_{i=0}^{n-1} \delta_i$, provided all $\delta_i \neq \perp$, and undefined, otherwise. We set $\overline{\delta}(E_P) := \sum_{i=0}^{n-1} |\delta_i|$ (which is always defined).

Given an edge vector sequence $(e_0, e_1, \ldots, e_{n-1})$, an edge vector e_i is called *inflection-edge vector* if $\delta_{i-1} \cdot \delta_i < 0$. In a polygon, the vertices preceding and succeeding an edge corresponding to an inflection-edge vector lie on different sides of the line along the edge. For example e_0 is an inflection-edge vector in Figure 3. An edge vector e_i is called *weak inflection-edge vector* if it belongs to a sequence of at least two edge vectors with the same edge angle bounded by turning angles with different sign, i.e. there exist $i_0 < i_1$ and $i_0 \leq i \leq i_1$ with $a_{i_0} = a_{i_0+1} = \cdots = a_i = \cdots = a_{i_1}$,

[1] e_i is *not* the edge (segment) connecting p_i and p_{i+1}, it is the vector from p_i to p_{i+1}.

and $\delta_{i_0-1} \cdot \delta_{i_1} < 0$. This means each weak inflection-edge vector e_i belongs to a sequence of weak inflection-edge vectors $e_{i_0}, e_{i_0+1}, \ldots, e_{i_1}$ and if the sequence of weak inflection-edge vectors were replaced by the sum of the weak inflection-edge vectors $e_{i_0} + e_{i_0+1} + \cdots + e_{i_1}$, the sum would be an inflection-edge vector.

Note that E_P determines P up to translation, and a sequence $(e_0, e_1, \ldots, e_{n-1})$ in $(\mathbb{R}^2 - o)^n$, $n \geq 2$, is the egde vector sequence of a polygon iff $\sum_{i=0}^{n-1} e_i = o$. We consider the values in Δ_P as real numbers and the arithmetic of these values without equivalence modulo 2π.

Observation 2.1 *Let P be a polygon with turning angles $(\delta_0, \delta_1, \ldots, \delta_{n-1})$, all $\delta_i \neq \perp$ (and hence $\delta(P)$ defined and $n \geq 3$).*

(i) $\delta(P)$ is a multiple of 2π.
(ii) If P is simple, then $\delta(P) \in \{2\pi, -2\pi\}$.
(iii) If P is convex, then either
  (a) $\delta(P) = 2\pi$ and $\delta_i \geq 0$ for all $i = 0, 1, \ldots, n-1$, or
  (b) $\delta(P) = -2\pi$ and $\delta_i \leq 0$ for all $i = 0, 1, \ldots, n-1$.
(iv) $\sum_{i=0}^{n-1} |\delta_i| \geq 2\pi$ with equality iff P is convex.

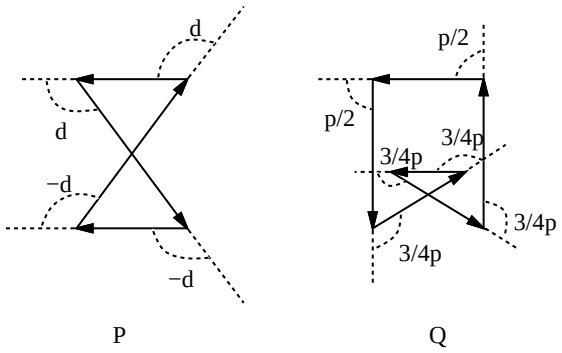

Figure 4: $\sum \delta_i = 0$ for P and $\sum \delta_i = 4\pi$ for Q.

Note that $\delta(P) = 2\pi$ and $\delta(P) = -2\pi$ discriminates whether we run through a simple polygon in counterclockwise or clockwise order, respectively. Without loss of generality, we assume that we run through a simple polygon in counterclockwise order.

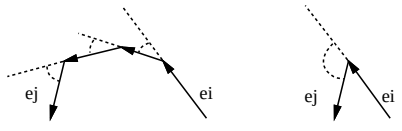

Figure 5: $\sum_{k=i}^{j-1} |\delta_k| = |\delta(e_i, e_j)|$.

Observation 2.2 *For a sequence of edge vectors* $(e_i, e_{i+1}, \ldots, e_j)$ *with* $\sum_{k=i}^{j-1} |\delta_k| < \pi$ *we have* $\sum_{k=i}^{j-1} |\delta_k| \geq |\delta(e_i, e_j)|$. *If, moreover,* $\delta_k \geq 0$ *for all* $k \in i, i+1, \ldots, j-1$ *or* $\delta_k \leq 0$ *for all* $k \in i, i+1, \ldots, j-1$, *then* $\sum_{k=i}^{j-1} |\delta_k| = |\delta(e_i, e_j)|$, *(see Figure 5).*

L_1-optimal merge. We call $Z = (z_0, z_1, \ldots, z_{n-1})$ an *L_1-optimal* cyclic merge of edge sequences X and Y if Z is a cyclic merge of X and Y and $\overline{\delta}(Z) = \sum_{i=0}^{n-1} |\delta(z_i, z_{i+1})|$ is minimal among all cyclic merges of X and Y.

Lemma 2.3 *For any cyclic merge Z of edge sequences X and Y, we have* $\max\{\overline{\delta}(X), \overline{\delta}(Y)\} \leq \overline{\delta}(Z)$.

Proof. Note that adding an edge into a sequence of edges cannot decrease its $\overline{\delta}$-value (recall Observation 2.2). Since we can obtain Z from X by successively adding the edges from Y, it follows that $\overline{\delta}(X) \leq \overline{\delta}(Z)$. Analogously, we can obtain Z starting from Y which gives $\overline{\delta}(Y) \leq \overline{\delta}(Z)$, and the lemma follows. □

With this lemma we are ready to prove the main property of L_1-optimal cyclic merges.

Lemma 2.4 *(i) If X and Y are edge sequences of convex polygons, then every L_1-optimal cyclic merge of X and Y is also convex.*
(ii) If X and Y are sequences of edge vectors of simple polygons and their sequences of
edge angles are cyclically equivalent (w.l.o.g. $x_i/||x_i|| = y_i/||y_i||$ for all $i \in 0, 1, \ldots n-1$), then
$Z = (x_0, y_0, x_1, y_1, \ldots x_{n-1}, y_{n-1})$ is an L_1-optimal cyclic merge of X and Y. Any L_1-optimal merge can be obtained from Z by successively swapping consecutive edge vectors e and e' with $e = \lambda e', \lambda > 0$.

Before we proceed with the proof, let us remark that (i) implies that condition C1 is satisfied. If X can be obtained from Y by scaling, then Z as described in (ii) corresponds to the cone section as required by condition C2. Since the swappings described do not change the actual surface (only its associated triangulation), this shows that C2 is also fulfilled.

Proof. (i) X comes from a convex polygon, if its angles are cyclically sorted. Two cyclically sorted sequences can be merged to a cyclically sorted sequence, which again describes a convex polygon. Since the $\overline{\delta}$-values of all these sequences are equal 2π, Lemma 2.3 or Observation 2.1 imply the claimed assertion. (ii) Since $\overline{\delta}(Z) = \overline{\delta}(X) = \overline{\delta}(Y)$, the optimality of Z follows immediately from Lemma 2.3. It remains to give the proof that any L_1-optimal merge can be obtained from Z by successively swapping consecutive edge vectors e and e' with $e = \lambda e', \lambda > 0$. This fact is somewhat more subtle, as it is perhaps witnessed by the fact that the statement becomes wrong, if we drop the assumption that X and Y come from simple polygons! As we will show at the end of the proof we can restrict ourselves to sequences of edge vectors without turning angles with value 0, i.e. $\delta(x_i, x_{i+1}) \neq 0$

for all $i \in 0, 1, \ldots, n-1$. If X is convex the claimed assertion follows directly from (i). Now we consider simple non-convex polygons X and Y. X must contain inflection-edge vectors otherwise all turning angles must be positive and since X is not convex, $\delta(X) > 2\pi$, a contradiction to X simple (Observation 2.1).

First observation: Since $\overline{\delta}(Z) = \overline{\delta}(X) = \overline{\delta}(Y)$ and Lemma 2.3, deleting z_j, $j \in 0, 1, \ldots, 2n-1$, from Z cannot decrease $\overline{\delta}(Z)$.

It follows that there is no inflection-edge vector in Z. It also follows that there is no undefined turning angle $\delta(z_j, z_{j+1})$ in Z. In this case deleting one of the edge vectors z_j or z_{j+1} would decrease $\overline{\delta}(Z)$. $|\delta(z_j, z_{j+1})| + |\delta(z_{j+1}, z_{j+2})| \geq \pi$ in $\overline{\delta}(Z)$ because $|\delta(z_j, z_{j+1})| = \pi$. Say z_{j+1} is deleted. Then $|\delta(z_j, z_{j+1})| + |\delta(z_{j+1}, z_{j+2})|$ is replaced by $|\delta(z_j, z_{j+2}|$ which is smaller than π. Since $\overline{\delta}(Z) = \overline{\delta}(X) = \overline{\delta}(Y)$, z_j and z_{j+2} are from different polygons and z_j and z_{j+1} are from different polygons and hence z_{j+1} and z_{j+2} are from the same polygon with the same turning angle. This contradicts to our assumption $\delta(x_i, x_{i+1}) \neq 0$ for all $i \in 0, 1, \ldots, n-1$ and therefore $|\delta(z_j, z_{j+2})|$ cannot be π.

Directly from the first observation follows a second observation: Let Z be an (I, J)-indexed cyclic merge of X and Y and $i < j$ be two consecutive indices in I or two consecutive indices in J. Then $z_i, z_{i+1}, \ldots, z_j$ is a convex sequence, i.e. $\delta(z_k, z_{k+1}) \geq 0$ for all $i \leq k \leq j-1$ or $\delta(z_k, z_{k+1}) \leq 0$ for all $i \leq k \leq j-1$ and $\sum_{k=i}^{j-1} \delta(z_k, z_{k+1}) = \delta(z_i, z_j)$ because from $\overline{\delta}(Z) = \overline{\delta}(X)$ we know that $\sum_{k=i}^{j-1} |\delta(z_k, z_{k+1})| = |\delta(z_i, z_j)|$.

Let $i_1 < i < i_2$ be consecutive in I and z_i corresponds to an inflection-edge vector in X. From the second observation it follows that $z_{i_1}, \ldots, z_i$ and $z_i, \ldots, z_{i_2}$ are convex sequences of edge vectors and $\delta(z_{i_1}, z_i) \cdot \delta(z_i, z_{i_2}) < 0$. Thus z_i must be a weak inflection-edge vector in Z because it cannot be an inflection-edge vector in Z (see above). This means that there is an adjacent edge vector, w.l.o.g. z_{i-1}, with index from J, $i-1 \in J$, has the same turning angle and is a weak inflection-edge vector in Z. Analogously, it corresponds to an inflection-edge vector in Y. The reason is if $j_1 < i-1 < j_2$ are consecutive in J then $z_{j_1}, \ldots, z_{i-1}$, overlapping with $z_{i_1}, \ldots, z_i$, and $z_{i-1}, \ldots, z_{j_2}$, overlapping with $z_i, \ldots, z_{i_2}$, are convex sequences of edge vectors and with $\delta(z_{i_1}, z_i) \cdot \delta(z_i, z_{i_2}) < 0$ also $\delta(z_{j_1}, z_{i-1}) \cdot \delta(z_{i-1}, z_{j_2}) < 0$. A sequence of weak inflection-edge vectors cannot be longer than 2 because of our restriction to sequences X without turning angles of value 0.

We argue analogously for each inflection-edge vector from Y. We conclude that there are no inflection-edge vectors, but pairs of weak inflection-edge vectors with the same turning angle, one from X and one from Y, and they correspond exactly to the inflection-edge vectors in X and Y.

Look at the sequence of inflection-edge vectors $(x_{i_0}, x_{i_1}, \ldots, x_{i_{k-1}})$ in X, $(y_{i_0}, y_{i_1}, \ldots, y_{i_{k-1}})$ in Y. If the pairs of weak inflection-edge vectors in Z are x_{i_l}, y_{i_l}, $0 \leq l \leq k-1$, then the proof can be completed using ideas analogously to (i) for the identical convex parts between x_{i_l} and $x_{i_{l+1}}$, and y_{i_l} and $y_{i_{l+1}}$.

Assume x_{i_l} is adjacent to $y_{i_{l+c}}$ for every l and a fixed integer constant c, $0 < c < k$. This means $\delta(x_{i_l-1}, x_{i_l}) = \sum_{k=r}^{j-1} \delta(z_k, z_{k+1}) = \delta(y_{i_{l+c}-1}, y_{i_{l+c}})$, this follows from

the fact that inflection-edge vectors from X and Y are exactly the weak inflection-edge vectors in Z and $\overline{\delta}(Z) = \overline{\delta}(X)$. Therefore $\delta(Z) = \delta(X) = 2\pi$. The sum of turning angles c times passing Z gives $c \cdot 2\pi$. This is

$$\begin{aligned} c \cdot \sum_{j=0}^{2n-1} \delta(z_j, z_{j+1}) &= \sum_{l=0}^{ck-1} \sum_{j=i_l}^{i_{l+1}-1} \delta(x_j, x_{j+1}) \\ &= \sum_{l=0}^{c-1} \sum_{j=i_l}^{i_{l+1}-1} \delta(x_j, x_{j+1}) + \sum_{l=c}^{2c-1} \sum_{j=i_l}^{i_{l+1}-1} \delta(x_j, x_{j+1}) + \cdots \\ &\quad + \sum_{l=k(c-1)}^{ck-1} \sum_{j=i_l}^{i_{l+1}-1} \delta(x_j, x_{j+1}) \\ &= k \cdot i \cdot 2\pi \end{aligned}$$

with $i \in \mathbb{Z}$ because the sum of turning angles between x_{i_l} and $x_{i_{l+c}}$ must be $i \cdot 2\pi$ since they have the same edge angle. We get $i = c/k$ which is a contradiction since $0 < c < k$ and i is an integer.

It remains to show how to handle sequences X with turning angles with value 0. Assume $\delta(x_{i_0}, x_{i_0+1}) = 0$ for some $i_0 \in 0, 1, \ldots, n-1$. A sequence of edge vectors X' is constructed by successively replacing pairs of consecutive edge vectors of X which have the same edge angle value ($x_{i_0}/||x_{i_0}|| = x_{i_0+1}/||x_{i_0+1}||$) by their sum $(x_{i_0} + x_{i_0+1})$, i.e. the corresponding sequence of edge angles is constructed by successively deleting edge angles where the preceding edge angle has the same value. Using Observation 2.2 and $\overline{\delta}(Z) = \overline{\delta}(X)$, it is easy to see that in an L_1-optimal cyclic merge of X and Y x_{i_0} and x_{i_0+1} may be consecutive or there are only edge vectors of Y with the same edge angle (as the one of x_{i_0} and x_{i_0+1}) in the sequence from x_{i_0} to x_{i_0+1}. It is not possible that in the sequence from x_{i_0} to x_{i_0+1} in the L_1-optimal cyclic merge of X and Y an edge angle with another value appears. So X can be reduced to X' and if the claim holds for X' then also for X. □

If the polygons are convex then the L_1-optimal cyclic merge corresponds to the Minkowski sum of the polygons. No such correspondence exists as soon as the polygons are not convex (e.g. the L_1-optimal merge is in general not unique).

L_2-optimal merge. A cyclic merge $Z = (z_0, z_1, \ldots, z_{m+n-1})$ of $X = (x_0, x_1, \ldots x_{n-1})$ and $Y = (y_0, y_1, \ldots y_{m-1})$ is *L_2-optimal* if $\sum_{i=0}^{m+n-1} \delta(z_i, z_{i+1})^2$ is minimal.

Lemma 2.5 *If X and Y are sequences of convex polygons then an L_2-optimal cyclic merge of X and Y is also convex.*

Proof. The proof of Lemma 2.5 is not as simple as the one for Lemma 2.4, because adding an edge vector to a sequence may actually decrease the sum of squares of turning angles.

Let $Z' = (z_0', z_1', \ldots, z_{n+m-1}')$ be a convex cyclic merge of X and Y, and $Z = (z_0, z_1, \ldots, z_{n+m-1})$ any non-convex cyclic merge of X and Y. We want to show

$$\sum_{j=0}^{m+n-1} \delta(z_j, z_{j+1})^2 > \sum_{j=0}^{m+n-1} \delta(z_j', z_{j+1}')^2.$$

We say that a (directed) interval z_a, z_b *covers* an interval z_c, z_d if the edge angles of z_c and z_d lie in the range of edge angles from z_a to z_b. Every interval z_j, z_{j+1}, $0 \leq j \leq n+m-1$, covers a non-empty sorted sequence of edge vectors of Z', i.e. $|\delta(z_j, z_{j+1})| = \sum_{i=k_j}^{l_j-1} \delta(z_i', z_{i+1}')$ and $z_j = z_{k_j}'$ and $z_{j+1} = z_{l_j}'$ or vice versa, $z_j = z_{l_j}'$ and $z_{j+1} = z_{k_j}'$. So we get

$$\sum_{j=0}^{n+m-1} \delta(z_j, z_{j+1})^2 = \sum_{j=0}^{n+m-1} (\sum_{i=k_j}^{l_j-1} \delta(z_i', z_{i+1}'))^2 \geq \sum_{j=0}^{n+m-1} \sum_{i=k_j}^{l_j-1} \delta(z_i', z_{i+1}')^2$$

If $\delta(Z) \geq 2\pi$ or $\delta(Z) \leq -2\pi$ there exists an index j for each interval z_i', z_{i+1}' with z_j, z_{j+1} covers z_i', z_{i+1}'.

$$\sum_{j=0}^{n+m-1} \sum_{i=k_j}^{l_j-1} \delta(z_i', z_{i+1}')^2 \geq \sum_{i=0}^{n+m-1} \delta(z_i', z_{i+1}')^2$$

The only possibility that equality holds is $k_j = l_j - 1$ for all j, $0 \leq j \leq m+n-1$, which means $Z = Z'$.
If $\delta(Z) = 0$ we cannot guarantee that each interval z_i', z_{i+1}' is covered by an interval z_j, z_{j+1}. But Z defines a polygon and therefore the union of all intervals z_j, z_{j+1}, $0 \leq j \leq n+m-1$, is connected and touches each z_j'. There is at most one index $i \in 0, 1, \ldots, n+m-1$, with interval z_i', z_{i+1}' is not covered by any interval z_j, z_{j+1}. The idea of the proof is to show that $\delta(z_i', z_{i+1}')^2$ is replaced by a larger value. All intervals building the subdivision of the remaining interval z_{i+1}', z_i' in Z' are covered by intervals in Z. They are already covered by positive intervals, i.e. pairs of consecutive edge vectors in Z with positive turning angle. We will show that there are additional costs of negative (which guarantees that they are additional) intervals in Z which are larger than $\delta(z_i', z_{i+1}')^2$. Two cases are possible: z_i', z_{i+1}' both come from one edge vector sequence, w.l.o.g. X, or one is from X and one from Y.

First case: $z_i', z_{i+1}' = x_j, x_{j+1}$, w.o.l.g. $j = 0$. We have to compare to $\delta(x_0, x_1)^2$ the additional costs of negative intervals of a solution omitting the interval x_0, x_1. Let y_0, y_1 be edge vectors of Y such that there is no edge vector of Y between y_0 and x_0 in Z and no edge vector of Y between x_1 and y_1 in Z, see Figure 6a). (In Figure 6, edge angles are displayed as points on the unit circle). $0 < \delta(x_0, x_1) < \pi$ and $0 < \delta(y_0, y_1) < \pi$ because X and Y are convex. In Z there must be edge vectors of Y between x_0 and x_1, $x_0 y_k y_{k+1} \ldots y_l x_1$, $l \geq k$, with $-\pi \leq \delta(x_0, y_k) < 0$ and $-\pi \leq \delta(y_l, x_1) < 0$, see Figure 6a). $\delta(x_0, y_k)^2 + \delta(y_l, x_1)^2 \geq \pi^2/2$ because

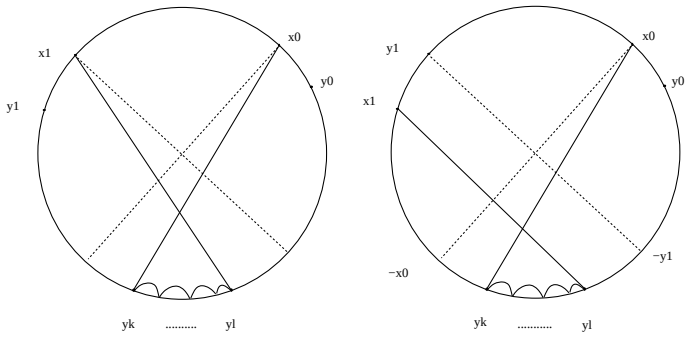

Figure 6: a) $z_i', z_{i+1}' = x_0, x_1$ b) $z_i', z_{i+1}' = x_0, y_1$

$\delta(x_0, y_k) + \delta(y_l, x_1) \leq -\pi$. Also there must be edge vectors of X between y_0 and y_1 in Z since Z does not cover the interval x_0, x_1. Assume the sequence of edge vectors of Z from x_0 to x_1 is part of the sequence from y_0 to y_1. This is a contradiction because in this case $y_k = y_1$, but $0 < \delta(y_0, y_1) < \pi$ and $0 > \delta(y_0, y_k) > -\pi$, or $y_l = y_0$ and $0 > \delta(y_l, y_1) > -\pi$. It follows that there is a sequence $y_0 x_{k'} \ldots x_{l'} y_1$ not containing x_0 or x_1 in Z with two negative intervals $y_0, x_{k'}$ and $x_{l'}, y_1$. $y_0, x_{k'}$ is negative because otherwise x_0 would be part of the sequence; $x_{l'}, y_1$ is negative because otherwise x_1 would be part of the sequence. For this sequence we also get additional costs of at least $\pi^2/2$. The additional costs sum up to at least $\pi^2 > \delta(x_0, x_1)^2$.

Second case: z_i', z_{i+1}' are from both polygons, one from X and one from Y, w.l.o.g. $z_i', z_{i+1}' = x_0, y_1$. In Z there must be edge vectors of Y between x_0 and x_1, $x_0 y_k y_{k+1} \ldots y_l x_1$, $l \geq k$, see Figure 6b). $\delta(x_0, y_k) < 0$ and $\delta(y_l, x_1) < 0$. Let $-x_0$ be the edge angle with direction of $x_0 - \pi$, $-y_1$ the edge angle with direction of $y_1 - \pi$. If $\delta(y_l, x_1)^2 > \delta(-x_0, -y_1)^2$ then Z is not L_2-optimal. Using this fact together with $0 > \delta(x_0, y_k) > -\pi$, we conclude that y_k, y_l is covered by $-x_0, -y_1$ and since $\delta(x_0, y_k) + \delta(y_l, x_1) \leq -\pi$, $\delta(x_0, y_k)^2 + \delta(y_l, x_1)^2 \geq \pi^2/2$. Analogously in Z there must be edge vectors of X between y_0 and y_1, $y_0, x_{k'}, \ldots, x_{l'}, y_1$, with interval $x_{k'}, x_{l'}$ covered by $-x_0, -y_1$ and analogously we get additional costs $\delta(y_0, x_{k'})^2 + \delta(x_{l'}, y_1) \geq \pi^2/2$. We conclude that Z cannot be an L_2-optimal cyclic merge. □

Condition C1 is obeyed, but L_2-optimal solutions may violate condition C2. To this end consider the example of two stars in Figure 7. Let the acute angle in the polygons be ϵ, $0 < \epsilon < \pi/2$. Then the L_2-value of the L_1 optimal solution is $4((\pi - \epsilon)^2 + (\pi/2 - \epsilon)^2)$. The alternative merge (given by a program as the L_2-optimal merge for $\epsilon \approx \pi/5$) has an L_2-value of $4(2(\pi/2 - \epsilon)^2 + (\pi/2 + \epsilon)^2 + \epsilon^2)$. As ϵ approaches 0, the first value converges to $5\pi^2$, while the second one converges to $3\pi^2$. So for some ϵ small enough ($\epsilon < (\sqrt{2} - 1)\pi/2$), the solution suggested by condition $C2$ will not be L_2-optimal.

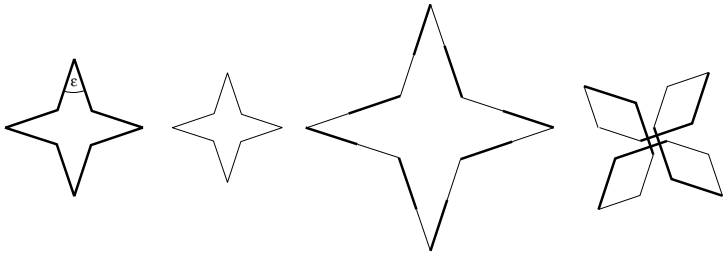

Figure 7: L_1-optimal merge and L_2-optimal merge of two cyclically equivalent polygons

3 Angle consistency

Let us right go back to the example in Figure 7. The solution suggested as L_2-optimal obviously leads to a surface with self-intersections, since the merged polygon Z is not simple; even without looking at the picture, we could compute $\delta(Z) = -4\pi$, a value which contradicts the simplicity of the underlying polygon (no matter what the lengths of the edge vectors are). In this section we will suggest a criterion which eliminates such obviously bad solutions.

Before we start with the key definition, we want to point out that a cyclic merge Z of two edge vector sequences X and Y does not necessarily determine the surface. However, the surface is determined if we give Z as an indexed merge[2], when it is clear which vector in Z comes from X and which one comes from Y.

Let $Z = (z_0, z_1, \ldots, z_{n-1})$ be an (I, J)-indexed merge of two edge vector sequences X and Y with $\delta(X)$ and $\delta(Y)$ defined. Let $i < j$ be two consecutive indices in I. We define $\delta_{i,j}^{(X)} := \delta(z_i, z_j)$, and $\delta_{i,j}^{(Z)} := \sum_{k=i}^{j-1} \delta(z_k, z_{k+1})$; analogously, we define $\delta_{i,j}^{(Y)}$ for consecutive indices in J. Moreover, we agree on the obvious cyclic extension for indices $i > j$, where i is the largest index in I and j is the smallest index in I (and similar for J).

We say that Z is *angle consistent*, if $\delta_{i,j}^{(X)} = \delta_{i,j}^{(Z)}$ and $\delta_{i,j}^{(Y)} = \delta_{i,j}^{(Z)}$ for all pairs of cyclically consecutive indices in I and J, respectively.

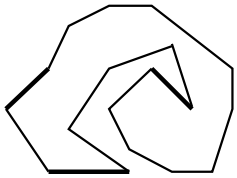

Figure 8: A cyclic merge which is simple but not angle consistent.

Note that if $0 \le i_0 \le i_1 \le \cdots i_{k-1} \le n-1$ are the indices in I, and X comes

[2]Recall definition in the beginning of Section 2.

from a simple polygon, then

$$2\pi = \delta(X) = \sum_{l=0}^{k-2} \delta(z_{i_l}, z_{i_{l+1}}) + \delta(z_{i_{k-1}}, z_{i_0}).$$

This sum equals $\delta(Z)$, if Z is an angle consistent merge; and hence $\delta(Z) = 2\pi$. However, it may very well be that $\delta(Z) = 2\pi$ (it may even be simple), but it is not angle consistent, see Figure 8.

Although, a cyclic merge which is not angle consistent may be simple, the resulting surface will always contain self-intersections (as will be shown below). So it is justified to exclude such merges for our surfaces. This will eliminate also self-intersections for our L_1-angle criterion; see Figure 9 for an example where an L_1-optimal merge violates angle consistency, and go back to Figure 2 for the L_1-optimal angle consistent merge.

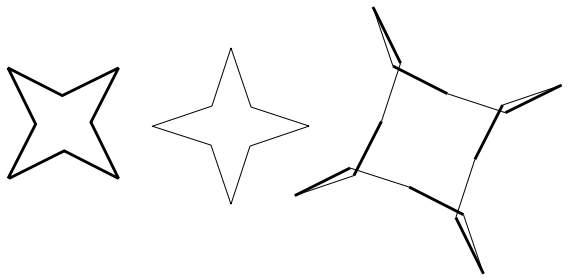

Figure 9: L_1-optimal cyclic merge which violates angle consistency.

Theorem 3.1 *An indexed cyclic merge of edge vector sequences of two simple polygons which is not angle consistent leads to a surface with self-intersections.*

Proof. Let us assume that P lies in the xy-plane, and Q lies in a parallel plane at height 1 (i.e. in the plane $z = 1$). We have argued before in the introduction, that the cyclic merge Z defined by a surface is the intersection of the surface with the plane at height $1/2$, scaled with a factor 2. Let us be more specific, saying that $Z = (z_0, z_1, \ldots, z_{n-1})$ is the (I, J)-indexed cyclic merge of the edge vector sequences of P and Q. If we consider now the intersection of the surface with a plane at height λ, $0 \leq \lambda \leq 1$, then this can be obtained from Z by multiplying all edge vectors from P (with index in I) by $1 - \lambda$, and the edge vectors from Q (with index in J) by λ. This gives a family of polygons with edge vector sequences Z_λ. The surface is free of self-intersections, if all polygons Z_λ are simple.

Assume Z contains an undefined turning angle and for this reason Z is not angle consistent. In this case no Z_λ is simple for $0 < \lambda < 1$. In the following we only consider sequences Z with $\delta(Z)$ defined.

For the remaining proof let us multiply the length of the edges in Z_λ, $0 < \lambda \leq 1$ by $1/\lambda$ to obtain edge vector sequences Z'_λ where the edge vectors from Q have constant length, and the edge vectors from P are multiplied by $(1 - \lambda)/\lambda$.

Consider now a violation of angle consistency, i.e. a pair $i < j$ of consecutive indices in I where $\delta_{i,j}^{(X)} \neq \delta_{i,j}^{(Z)}$ (other cases of violation are symmetric). Hence, the sequences Z'_λ contain as a subsequence $S_\mu = (\mu z_i, z_{i+1}, \ldots, z_{j-1}, \mu z_j)$ with $\mu := (1-\lambda)/\lambda \longrightarrow \infty$ as $\lambda \longrightarrow 0$.

If $0 < \delta_{i,j}^{(X)} < \pi$, let $v_\mu = -(\mu z_i + z_{i+1} + \cdots + z_{j-1} + \mu z_j)$, i.e. $(\mu z_i, z_{i+1}, \ldots, z_{j-1}, \mu z_j, v_\mu)$ is the edge vector sequence of a polygon, unless $v_\mu = o$; if $v_\mu = o$, then this immediately reveals a self-intersection. Let $\mu \longrightarrow \infty$, v_μ only intersects μz_i and μz_j in the polygon. If the polygon is simple then $\delta_{i,j}^{(X)} = \delta_{i,j}^{(Z)}$ which is a contradiction to the assumption. Otherwise there is a part of Z_λ which is not simple and which yields a self-intersection. In a similar way the case $\delta_{i,j}^{(X)} = 0$ can be handled (let $v_\mu^1 + v_\mu^2 = -(\mu z_i + z_{i+1} + \cdots + z_{j-1} + \mu z_j)$ and $\delta(v_\mu^1, v_\mu^2) = \epsilon \longrightarrow 0$.) □

If $X = (x_0, x_1, \ldots, x_{n-1})$ and $Y = (y_0, y_1, \ldots, y_{m-1})$ are edge vector sequences of simple polygons then there always exists a cyclic merge $Z = (x_0, \ldots, x_i, y_0, y_1, \ldots, y_{m-1}, x_{i+1}, \ldots, x_{n-1})$ of X and Y which is angle consistent. For example take the leftmost vertex of X (or the uppermost of these if there are more than one) as the i-th vertex and the rightmost (the lowermost of those) of Y as the $(m-1)$-th then $\delta(Z) = 2\pi$ and it directly follows from the construction that angle consistency is fulfilled.

Angle consistency does not concern conditions C1 and C2. If an optimal cyclic merge fulfills the conditions then also the optimal among the angle consistent fulfills the conditions.

4 Algorithm

If one polygon is convex it is easy to find an L_1-optimal cyclic merge.

Lemma 4.1 *An angle consistent L_1-optimal cyclic merge Z of an edge sequence X of a convex polygon with n vertices and an edge sequence Y of a simple polygon with m vertices can be constructed in $O(n+m)$ time.*

Proof. The angles of X are cyclically sorted. Edges of X can be successively inserted into the edge sequence of Y without increasing its $\bar{\delta}$-value because $\delta(Y) = 2\pi$. If this is done in a greedy way (insert as soon as possible), angle consistency is guaranteed. □

If none of the polygons is convex the problem can be formulated as a shortest path problem in a directed graph.

Description of the algorithm:
Every possible triangle in a connecting surface (defined by an edge of one polygon and a vertex of the other) is represented by a node in the graph. The node set of the graph has cardinality $2 \cdot m \cdot n$. A node is labeled $(i, j, 0)$ if the triangle is defined as the convex hull of the edge between the $(i-1)$-th and i-th vertex of

polygon P and the j-th vertex of polygon Q; $(i,j,1)$ is defined analogously by the i-th vertex of P and the $j-1$-th and j-th of Q. Arcs in the graph connect nodes of consecutive triangles which share an edge connecting P and Q. The graph is a torus graph. Indegree and outdegree of a vertex are 2. Arc weights are assigned according to the absolute value of the turning angle between the polygon edges of the two consecutive triangles. Fixing a starting triangle (w.l.o.g. $(0,j,\cdot)$), we are looking for a cycle of minimum weight passing node $(0,j,\cdot)$ containing $n+m$ triangles. A global optimal solution is the minimum among all minimum weight cycles in the torus graph passing $(0,0,0),(0,0,1),(0,1,0),(0,1,1),(0,2,0),\ldots,(0,m-1,0)$ or $(0,m-1,1)$, respectively. For a fixed starting triangle, w.l.o.g. $(0,0,0)$, we regard a subgraph of the torus graph which is a directed acyclic graph with $2\cdot(n+1)\cdot(m+1)$ nodes $(i,j,0)$ and $(i,j,1)$; $0 \le i \le n$, $0 \le j \le m$, where $(n,m,\cdot)$ is a copy of $(0,0,\cdot)$. A minimum weight cycle in the torus graph passing $(0,0,0)$ corresponds to a shortest path from $(0,0,0)$ to $(n,m,0)$ in the subgraph. A shortest path can be computed in $O(n\cdot m)$ time since the subgraph is a directed acyclic graph of this size. But we have to compute a shortest path for each of the $2\cdot m$ starting triangles $(0,0,0),(0,0,1),(0,1,0),\ldots,(0,m-1,1)$. So the overall running time to compute the value of an L_1-optimal cyclic merge is $O(n\cdot m^2)$ and space requirements are $O(n\cdot m)$ (the number of nodes of the union of the subgraphs is $2\cdot(n+1)(2m)$). The L_1-optimal merge itself can be obtained by backtracking through the graph.

Theorem 4.2 *An L_1-optimal merge of two polygons with n and m vertices can be computed in $O(n\cdot m^2)$ time.*

The algorithm can be used to compute other angle dependent optimal merges like the L_2-optimal merge.

Remark. This solution is based on two papers, one of the first papers written on contour triangulation [K], it employs a smaller directed graph to compute a maximal volume contour triangulation; Fuchs, Kedem and Uselton [FKU] refined the modeling of the graph to accelerate the algorithm. They gave a faster algorithm with running time $O(n\cdot m\cdot\log m)$ but they need graph planarity and our subgraphs are not planar. Sloan and Painter [SP] also used this approach and suggested a heuristic to improve the graph search.

The L_1-optimal merge produced by the algorithm may not fulfill angle consistency. To guarantee that the solution is angle consistent we have to extend the algorithm.

Suppose starting vertex $(0,0,0)$ is fixed. (We proceed analogously for all $2m$ starting vertices.) Guaranteeing angle consistency, the algorithm successively computes shortest paths to all vertices of the graph. Reaching a vertex we test if the path represents an angle consistent part of a solution. For example if the vertex corresponds to a triangle with a polygon edge from edge vector z_j with j in I and $i<j$ consecutive indices in I, we test if $\delta^{(X)}_{i,j}=\delta^{(Z)}_{i,j}$. To do this test in constant time per vertex we compute two entries δ^X and δ^Y at every vertex. Reaching a vertex corresponding z_j, δ^X denotes $\delta^{(Z)}_{k,j}$ with k is the largest index in I smaller than j, and δ^Y denotes $\delta^{(Z)}_{l,j}$ with l the largest vertex smaller j in J. δ^X and δ^Y are computed

and updated in constant time per vertex. If j is from I then $\delta^{(Z)}_{i,j}$ is given by δ^X. If j is from J and $i < j$ preceding j in J then we have to test if $\delta^{(Y)}_{i,j} = \delta^{(Z)}_{i,j}$ and $\delta^{(Z)}_{i,j}$ is given by δ^Y.

For every vertex $(i,j,0)$ $[(i,j,1)]$ in the graph the shortest angle consistent path from $(0,0,0)$ passing $(i,j,0)$ $[(i,j,1)]$ to $(i+1,j,0)$ and $(i,j+1,1)$ is computed. This means that we check what will happen if the next edge vector from X or from Y is taken. It is easy to compute these angle consistent paths for $i = 0$ or $j = 0$ (assuming $\delta^{(Y)}_{k,l} = \delta^{(Z)}_{k,l}$ with $z_k = y_0$ and $z_l = y_1$). Now we compute the paths to $(i,j,.)$ vertex by vertex in rows, what means before j is increased all values for $(i,j,.)$ for all $0 < i \leq n$ are computed. At every vertex the shortest angle consistent paths are computed as the shortest paths in the algorithm above; only if the shortest angle consistent paths passing $(i,j,0)$ $[(i,j,1)]$ to $(i,j+1,1)$ $[(i+1,j,0)]$ are computed angle consistency may be violated. Suppose angle consistency is violated at $(i,j,0)$, i.e. passing $(i,j,0)$ taking the arc to $(i,j+1,1)$. Only the pair of edge vectors $y_j(= z_k)$ and $y_{j+1}(= z_l)$ violates angle consistency, $\delta^{(Y)}_{k,l} \neq \delta^{(Z)}_{k,l}$. The shortest angle consistent path we are looking for contains a shortest angle consistent subpath passing $(i',j,1)$ taking the arc to $(i'+1,j,0)$ which we already computed for all $i' < i$. For each i' compute the length of the path passing $(i',j,1)$, $(i'+1,j,0),\ldots,(i,j,0),(i,j+1,1)$ and test if $\delta^{(Y)}_{k',l} = \delta^{(Z)}_{k',l}$ where Z is the merge corresponding to the path and $x_{i'} = z_{k'}$. All together for all i' this can be done in $O(n)$ time and also finding the shortest angle consistent among these $O(n)$ paths takes the same time. With this algorithm we find the shortest angle consistent path, we guaranteed angle consistency for all pairs of consecutive indices in I and J but k,l with $z_k = y_0$ and $z_l = y_1$ (see above). But $\delta^{(Y)}_{k,l} = \delta^{(Z)}_{k,l}$, because $\delta(X) = \delta(Z) = 2\pi = \sum \delta^{(Z)}_{i,j}$ with summation over all pairs i,j of consecutive indices in J.

At each vertex we spend at most $O(n)$ time. The resulting running time for a fixed starting vertex is $O(n^2 \cdot m)$ time, the overall running time to compute an L_1-optimal angle consistent merge, i.e. L_1-optimal among the angle consistent, is $O(n^2 \cdot m^2)$ time.

For many polygons it is possible to compute an L_1-optimal angle consistent merge in less time. We exploit the degree of convexity of a polygon in a similar way to Lemma 4.1. Given an edge vector sequence $X = (x_0, x_1, \ldots, x_{n-1})$, we define

$$d(X) = \max_{i,j}\{\textstyle\sum_{k=i}^{j-1} |\delta_k|| \text{ with } \quad \delta_k \geq 0 \text{ for all } k \in i, i+1, \ldots, j-1 \text{ or } \\ \delta_k \leq 0 \text{ for all } k \in i, i+1, \ldots, j-1\}$$

The *distortion* d_X of X is defined as $d_X := \lfloor d(X)/\pi \rfloor$. (This is a notion related e.g. to the winding number in [GRS].) Recall the definition of an inflection-edge vector in the beginning of Section 2. The number of inflection-edge vectors and sequences of weak inflection-edge vectors describes the degree of "convexity" of an edge vector sequence and the distortion describes how "spiral" it is.

Theorem 4.3 *Let X and Y be edge vector sequences of two simple polygons with n and m points and d is the maximum of their distortion. t is the number of inflection-edge vectors plus the number of sequences of weak inflection-edge vectors of X and*

Y. Then an L_1-optimal angle consistent cyclic merge of X and Y can be constructed in $O((dt)^4 + n + m)$ time.

Proof. $X = (x_0, x_1, \ldots, x_{n-1})$ and $Y = (y_0, y_1, \ldots, y_{m-1})$ are the edge vector sequences of two simple polygons. If X or Y is an edge vector sequence of a convex polygon then Lemma 4.1 proves the statement of this theorem. Similarly to Lemma 2.4 we assume that X and Y do not contain turning angles with value 0. Therefore we also assume that t is the number of inflection-edge vectors. The reason is given at the end of the proof.

X decomposes into maximal convex chains, i.e. subsequences $(x_i, x_{i+1}, \ldots, x_j)$ with $\delta_k > 0$ for all $k \in \{i, i+1, \ldots, j-1\}$ with $\delta_{i-1} < 0$ and $\delta_{j+1} < 0$, or $\delta_k < 0$ for all $k \in \{i, i+1, \ldots, j-1\}$ with $\delta_{i-1} > 0$ and $\delta_{j+1} > 0$. (Y analogously). Notice that the number of maximal convex chains in X and Y is t.

We will proceed as follows: First X and Y are reduced to at most dt edge vectors. Then a partial solution for the reduced problem is computed with the algorithm above in $O((dt)^4)$ time. In the second step the removed edge vectors are merged into the partial solution in $O(n+m)$ time and we get an L_1-optimal angle consistent cyclic merge of X and Y.

Reduction of X to X' and Y to Y': X' contains all inflection-edge vectors of X, these are the first and last edge vectors of the maximal convex chains, together with some additional edge vectors which witness the spirals of the polygon. Suppose x_i and x_j are the first and last edge vector of a positive maximal convex chain (consecutive inflection-edge vectors) and $\sum_{k=i}^{j-1} \delta(x_k, x_{k+1}) > \pi$. A negative maximal convex chain will be handled analogously. Beginning with $x_i = x_{i_0}$ (walking in direction x_j) we take from X the next possible edge vector x_{i_1} with $\sum_{k=i_0}^{i_1-1} \delta(x_k, x_{k+1}) < \pi$ and $\sum_{k=i_0}^{i_1} \delta(x_k, x_{k+1}) > \pi$. If $\sum_{k=i_1}^{j-1} \delta(x_k, x_{k+1}) > \pi$ then beginning with x_{i_1}, we take the last possible edge vector x_{i_2} with $\sum_{k=i_1}^{i_2-1} \delta(x_k, x_{k+1}) < \pi$ etc. until we have taken x_{i_l} with $\sum_{k=i_l}^{j-1} \delta_x(x_k, x_{k+1}) < \pi$. ($Y'$ analogously). In X' and Y' we have added at most d edge vectors per inflection-edge vector. With the above algorithm an L_1-optimal angle consistent cyclic merge Z' of X' and Y' is computed. Assume Z is the L_1-optimal angle consistent cyclic merge of X and Y. $\overline{\delta}(Z') \leq \overline{\delta}(Z)$ because $X' \subseteq X$ and $Y' \subseteq Y$. More precisely: Let Z'' be generated from Z by deleting the edge vectors lying in X but not in X' and those lying in Y but not in Y'. $\overline{\delta}(Z'') \leq \overline{\delta}(Z)$ (Observation 2.2) and $\overline{\delta}(Z') \leq \overline{\delta}(Z'')$ because Z' is the L_1-optimal angle consistent cyclic merge of X' and Y' and Z'' is an angle consistent cyclic merge of the same edge vectors.

Merging step: The edges removed from X, $X - X'$, consist of sorted sequences which are merged into Z' in a way described in Lemma 4.1 such that the ordering of the edge vectors relative to X is preserved. For the resulting cyclic merge Z'_X it holds: $\overline{\delta}(Z'_X) = \overline{\delta}(Z')$. The edge vectors of $Y - Y'$ are merged into Z'_X in the same way and we get Z'_{XY}. $\overline{\delta}(Z'_{XY}) = \overline{\delta}(Z') \leq \overline{\delta}(Z)$ and also $\overline{\delta}(Z) \leq \overline{\delta}(Z'_{XY})$ because of the optimality of Z. It follows $\overline{\delta}(Z) = \overline{\delta}(Z'_{XY})$.

It remains to show why we can restrict ourselves to sequences X and Y without turning angles of value 0. For each sequence of consecutive turning angles with

value 0 we replace the edge vectors defining those turning angles by their sum. The modified sequences $\tilde{X}$ and $\tilde{Y}$ do not contain adjacent edge vectors with the same edge angle. Each sequence of weak inflection-edge vectors becomes an inflection-edge vector in the modified sequence, i.e. t remains the same. The L_1-optimal solution $\tilde{Z}$ of the modified sequences induces an L_1-optimal solution Z of X and Y by backwards replacing the sums of edge vectors by the corresponding sequences. Z is an L_1-optimal cyclic merge for X and Y since the value of the solution, $\overline{\delta}(Z)$, remains the same as for the modified sequences, $\overline{\delta}(\tilde{Z})$. $\overline{\delta}(Z) \geq \overline{\delta}(\tilde{Z})$ since the sequence of edge angles of $\tilde{X}$ resp. $\tilde{Y}$ is a subsequence of edge angles of X resp. Y and restricting Z to $\tilde{X}$ and $\tilde{Y}$ cannot give a better value than $\overline{\delta}(\hat{Z})$ (Observation 2.2) . □

5 Experimental results

Figure 10: Synthetic example.

We have implemented the algorithm on a SUN Sparc 10 in C. The software of the algorithm consists of about 3000 lines of code and additional 6000 lines of code which contains an editor for creating synthetic examples, support for the graphics output and additional code for other optimality criteria. To compare the constructed surfaces we have implemented four optimality criteria. An area-optimal surface, L_1-optimal and L_2-optimal solutions, and a 'smoothest' surface can be computed. A smoothest surface is a surface where the sum of the absolute values of angles between normal vectors of consecutive triangles is minimized. A similar criterion is used in the context of reconstructing surfaces from a given set of points in $\mathbb{R}^3$ [ChShYL], [DLR]. Although considering the smoothest surface is intuitively appealing neither condition C1 nor condition C2 can be guaranteed by the smoothest surface.

To get an impression of the performance and characteristics of the algorithm using angle criteria we present some specific examples, in the beginning two synthetic examples to demonstrate the characteristics of L_1-optimal solutions:

Figure 10 represents the top view of two oval contours which have to be connected by a surface. An adequate solution is given by the L_1-optimal merge. Since the polygons are convex the connecting surface corresponding to the L_1- or L_2-optimal merge is convex. In comparison the area-optimal solution is shown in the right part of the Figure.

Figure 11: Synthetic example.

Figure 11 shows an L_1-optimal and an area-optimal solution of two rectangles with 'peaks' on the same side but at a different position. Since the L_1-optimal solution does not depend on edge lengths (but on edge angles only) the two peaks are connected. Some of the triangles of the resulting surface are slanted. (We have a natural example of a face where this leads to undesired effects.) The area-optimal solution depends very much on the position of the two polygons, it connects the peaks to the nearest point in the other polygon. Without knowing the application it is difficult to decide which solution is the better one. It could be desirable to connect special similar features of the polygons, for instance if the cross-sections represent 2D animation.

The following examples are results from the execution of the algorithm on medical data.

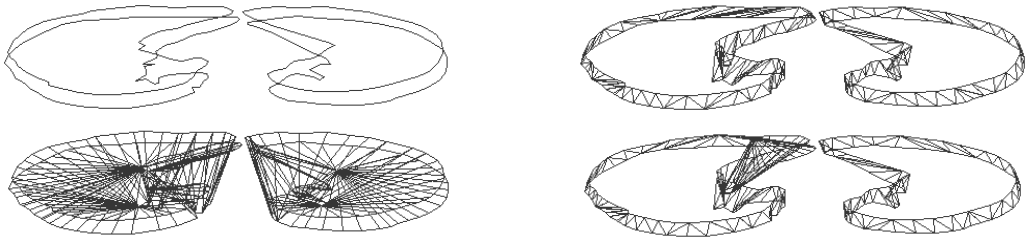

Figure 12: Cross-sections from the lungs.

Figure 12 shows two consecutive cross-sections from the lungs. In each cross-section two polygons are displayed, the two lobes of the lungs. The L_1-optimal is depicted in the upper right part of Figure 12 and the L_2-optimal solution in the lower left. While the L_1-optimal surface represents an adequate solution, the L_2-optimal and the 'smoothest' surface which is similar to the L_2-optimal are twisted surfaces where large portions of one polygon are connected to one point of the other. In the lower right part of Figure 12 we can observe the effect of adding the angle consistency condition, it shows the L_2-optimal among the angle-consistent. In the right lobe of the lungs there remains no self-intersection and the surface is 'intuitively correct'. Although the solution is angle consistent, in general we cannot guarantee that there is no self-intersection; consider the left lobe of the lungs, we see that the

sharp turning angle of the lower polygon is connected to a sequence of edges of the upper polygon because the value of the sharp turning angle in the merge is reduced by inserting edges of the other polygon.

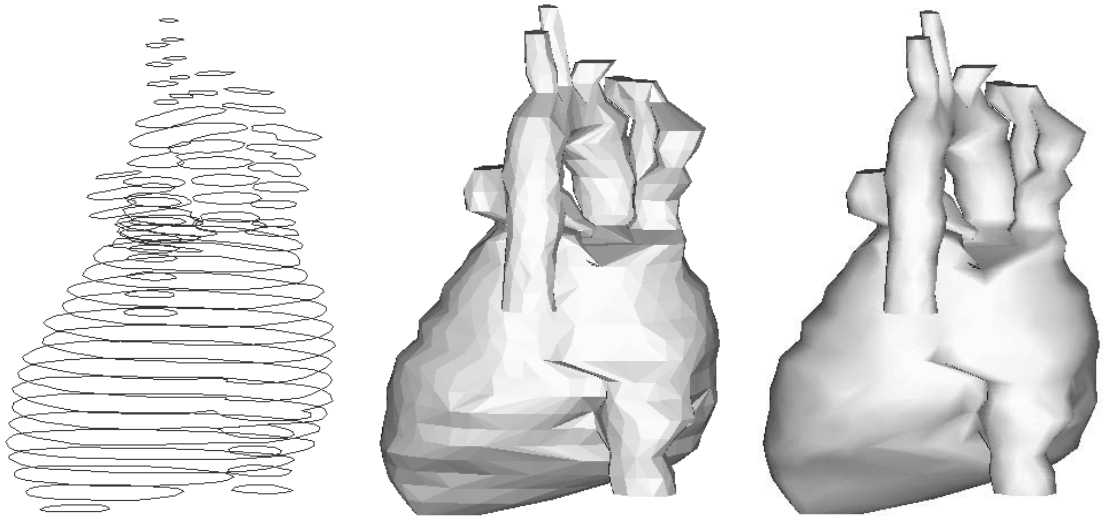

Figure 13: Reconstructed heart.

The next example, Figure 13 shows the set of contours of a heart and a shaded and a Gouraud-shaded display of the reconstruction.

	cross-sections	points	contours	CPU time
heart	30	1280	65	2.6 s
lungs	34	3121	88	22.4 s
hip	34	1739	39	11.2 s
head	17	856	26	6.2 s

Table 1: Some experimental results.

Table 1 sums up the running time for computing the L_1-optimal angle consistent solutions of some experiments. We observe that the running time of course depends on the number of points and number of contours but most important is the shape, i.e. the degree of convexity of the cross-sections. For example the heart consists of large convex parts whereas the data of the head contain many concavities and the running time for the reconstruction of the heart is less than half of time for reconstructing the head although there are 50% more points.

References

[BS] G. Barequet, M. Sharir. Piecewise-linear interpolation between polygonal slices. *Technical Report 275/93 Eskenasy Institute of Computer Science School of Mathematical Sciences, Tel-Aviv University,* 1993

[BGLS] C. Barillot, B. Gilbaud, L. M. Luo, J. M. Scarabin. 3D representation of anatomic structures from CT examination. *Biostereometrics*, 1985, 307-314

[B] J. D. Boissonnat. Shape reconstruction from planar cross-sections. *Computer Vision, Graphics and Image Processing 44*, 1988, 1-29

[BG] J. D. Boissonnat, B. Geiger. Three dimensional reconstruction of complex shapes based on the Delaunay triangulation. *Technical Report 1697, INRIA Sophia Antipolis,* 1992

[ChShYL] B. K. Choi, H. Y. Shin, Y. I. Yoon, J. W. Lee. Triangulation of scattered data in 3D space. *Computer Aided Design 20,* 1988, 239-248

[ChrS] H. N. Christiansen, T. W. Sederberg. Conversion of complex contour line definitions into polygonal element mosaics. *Computer Graphics 13,* 1978, 187-192

[C] P. N. Cook et al. Three-dimensional reconstruction from cross-sections for medical applications. *Proceedings, 14th Hawaii Int. Conf. on System Sci.,* 1981, 358-389

[DLR] N. Dyn, D. Levin, S. Rippa. Data dependent triangulations for piecewise linear interpolation. *IMA Journal of Numerical Analysis 10,* 1990, 137-154

[FKU] H. Fuchs, Z. M. Kedem, S. P. Uselton. Optimal surface reconstruction from planar contours. *Communications of the ACM 20,* 1977, 693-702

[GD] S. Ganapathy, T. G. Dennehy. A new general triangulation method for planar contours. *ACM Trans. Computer Graphics 16,* 1982, 69-75

[GORS] C. Gitlin, J. O'Rourke, V. Subramanian. On Reconstructing Polyhedra from Parallel Slices. *Technical Report 025, Smith College, Dept. Computer Science,* 1993

[GRS] L. Guibas, L. Ramshaw, J. Stolfi. A kinetic framework for computational geometry. *Proc. 24th IEEE Foundations of Computer Science,* 1983, 100-111

[K] E. Keppel. Approximating complex surfaces by triangulation of contour lines, *IBM Journal of Research and Development 19,* 1975, 2-11

[LC] W. Lorensen, H. Cline. Marching cubes: a high resolution 3D surface construction algorithm. *Computer Graphics 21,* 1987, 163-169

[MK] H. Müller, A. Klingert. Surface interpolation from cross sections. *in: H. Hagen, H. Müller, G.M. Nielson (eds.), FOCUS ON SCIENTIFIC VISUALIZATION,* 1993, 139-189

[MSS] D. Meyers, S. Skinner, K. Sloan. Surfaces from Contours. *ACM Transactions on Graphics, Vol. 11, No. 3,* 1992, 228-258

[Sch] L. L. Schumaker. Reconstructing 3D objects from cross-sections, *in: W. Dahmen et. al. (eds.), COMPUTATION OF CURVES AND SURFACES,* 1990, 275-309

[SG] T. W. Sederberg, E. Greenwood. A physically based approach to 2-D shape blending. *SIGGRAPH '92 Conference Proceedings, Computer Graphics, Vol. 26, No. 2,* 1992, 25-34

[SP] K. R. Sloan, J. Painter. Pessimal guesses may be optimal: A counterintuitive search result. *IEEE Transactions on Pattern Analysis and Machine Intelligence 10,* 1988, 949-955

[WA] Y. F. Wang, J. K. Aggarwal. Surface reconstruction and representation of 3D scenes. *Pattern Recognition 19,* 1986, 197-207