

ON A GENERALIZATION OF SCHÖNHARDT’S POLYHEDRON

JÖRG RAMBAU

ABSTRACT. We show that the non-convex twisted prism over an n -gon cannot be triangulated without new vertices. For this, it does not matter what the coordinates of the n -gon are as long as the top and the bottom n -gon are congruent and the twist is not too large. This generalizes Schönhardt’s polyhedron, which is the non-convex twisted prism over a triangle.

1. THE BACKGROUND

Back in 1911, Lennes [4] presented the first simple three-dimensional non-convex polyhedron whose interior cannot be triangulated without new vertices. The more famous example, however, was given in 1927 by Schönhardt: he observed that in the non-convex twisted triangular prism (subsequently called “Schönhardt’s polyhedron”) every diagonal that is not a face lies completely in the exterior [6]. This implies immediately that there can be no triangulation of it without new vertices because there is simply no interior tetrahedron: all possible tetrahedra spanned by four of its six vertices would introduce new edges. Moreover, he proved that every simple polyhedron with the same properties must have at least six vertices. Later, further such non-convex, non-triangulable polyhedra with an arbitrary number of points have been presented. Among them, Bagemihl’s polyhedron [1] also has the feature that every non-facial diagonal is in the exterior.

The twisted prism over an arbitrary n -gon would arguably be the most natural generalization of Schönhardt’s polyhedron. Surprisingly enough, there has been no proof so far that it cannot be triangulated without new vertices. One of the reasons seems to be that—in contrast to Schönhardt’s and Bagemihl’s polyhedra—not every non-facial diagonal lies completely outside the polygonal prism. Yet, the non-convex twisted polygonal prism can indeed not be triangulated without new vertices, as we will show below. For this, it does not matter what the coordinates of the n -gon are as long as the top and the bottom n -gon are congruent and the twist is just a perturbation by rotation, i.e., it is not too large.

There is a relation between Schönhardt’s polyhedron and the (untwisted) triangular prism with a prescribed boundary triangulation that is cyclically symmetric: there is no triangulation of the triangular prism that extends a cyclically symmetric boundary triangulation without new vertices. Similarly, there is no a triangulation of the general polygonal prism inducing a cyclically symmetric triangulation of the boundary quadrilaterals.

Besides the fact that the (frequently asked) question about the existence of triangulations of the non-convex twisted polygonal prism deserves a conclusive answer at last, we

Date: September 25, 2003.

Supported by a general membership at the Mathematical Sciences Research Institute Berkeley (MSRI) in the special semester “Discrete and Computational Geometry” (2003); research at MSRI is supported in part by NSF grant DMS-9810361.

mention one other motivation for studying problems like this. Deciding the existence of a triangulation without new vertices for a fixed polyhedron is NP-hard [5]. In studying the twisted polygonal prism we surprisingly hit the borderline between existence and non-existence of triangulations without new vertices in a single type of point configurations, and this could make the twisted or untwisted polygonal prism a handy gadget for NP-hardness proofs.

2. THE OBJECTS

Consider a two-dimensional point configuration $C_n := \{v_0, v_1, \dots, v_{n-1}\}$ in strictly convex position labeled counter-clockwise. Fix a point o in the interior of C_n in \mathbb{R}^2 . For $\alpha \in [0, 2\pi)$, let $C_n(\alpha)$ be a copy of C_n rotated by α around the point o (rotation by an angle in $(0, 2\pi)$ means counter-clockwise rotation). We call the corresponding points w_0, w_1, \dots, w_{n-1} . The *Cayley embedding* of C_n and $C_n(\alpha)$ is defined by

$$P_n(\alpha) := \text{conv}((C_n \times \{0\}) \cup (C_n(\alpha) \times \{1\})).$$

A *triangulation* of a three-dimensional polyhedron P is a dissection into finitely many tetrahedra such that any two intersect in a common face (possibly empty). For a triangulation of P and a simplex F of arbitrary dimension we say T *uses* F if F is a face of some tetrahedron in T . Faces are denoted by their sets of vertices. A *triangulation without new vertices* or a *v -triangulation* of P is a triangulation all of whose vertices are vertices of P .

$P_n := P_n(0)$ is known as a *prism* over C_n . The *cyclic set of diagonals*

$$D_c := \{\{v_i, w_{i+1}\} : i = 0, 1, \dots, n-1\}$$

induces a triangulation of the quadrilateral facets of $P_n(0)$ into the triangles $\{v_i, w_i, w_{i+1}\}$ and $\{v_i, w_{i+1}, v_{i+1}\}$, $i = 0, 1, \dots, n-1$ (all indices regarded modulo n).

The continuity of the determinant function ensures that there is an $\alpha > 0$ such that no full-dimensional tetrahedron in $P_n(0)$ has a reversed orientation (sign of determinant of the points in homogeneous coordinates) in $P_n(\alpha)$. In that case, the *vertical edges* $\{v_i, w_i\}$ and the *reverse cyclic edges* $\{w_i, v_{i+1}\}$ are among the boundary edges of $P_n(\alpha)$, for all $i = 0, 1, \dots, n-1$. For such an α , we call $P_n(\alpha)$ a *convex twisted prism* over C_n . ($P_n(\alpha)$ is a convex twisted prism over C_n if and only if the map sending $v_i, w_i \in P_n(\alpha)$ to the corresponding $v_i, w_i \in P_n(0)$ induces a weak map of oriented matroids [2].)

For a convex twisted prism over C_n , the *cyclic set of tetrahedra* is the set of tetrahedra

$$T_c := \{\{v_i, v_{i+1}, w_i, w_{i+1}\} : i = 0, 1, \dots, n-1\}.$$

Any two of these tetrahedra intersect in a common edge.

3. THE RESULTS

Theorem 1. *For all $n \geq 3$, no prism $P_n(0)$ over an n -gon admits a triangulation without new vertices that uses the cyclic set D_c of diagonals.*

Theorem 2. *For all $n \geq 3$, no convex twisted prism $P_n(\alpha)$ admits a triangulation that contains the cyclic set T_c of tetrahedra.*

We define the *non-convex twisted prism* $\check{P}_n(\alpha)$ to be the topological closure of $P_n(\alpha) \setminus T_c$. Since the twist is not too large, this is a non-convex simple polyhedron. Here is now the generalization of Schönhardt's polyhedron:

Corollary 1. *For all $n \geq 3$, the non-convex twisted prism $\check{P}_n(\alpha)$ cannot be triangulated without new vertices.*

Remark 1. When C_n is a regular triangle and $\alpha \in (0, 2\pi/3)$, the twisted prism $P_3(\alpha)$ coincides with Schönhardt's twisted prism.

4. THE TOOLS

For a more detailed background about the following consult [3] and the references therein.

4.1. Minkowski sums and mixed subdivisions. Let P and Q be point configurations in \mathbb{R}^2 . Then the *Minkowski sum of P and Q scaled by $\lambda \in (0, 1)$* is the point configuration

$$(1 - \lambda)P + \lambda Q := \{(1 - \lambda)p + \lambda q : p \in P, q \in Q\} \subset \mathbb{R}^2.$$

We make the following simplifying assumption: we consider only generic $\lambda \in (0, 1)$, for which $(1 - \lambda)p + \lambda q = (1 - \lambda)p' + \lambda q'$ implies that $p = p'$ and $q = q'$. A *mixed cell* in $(1 - \lambda)P + \lambda Q$ is the Minkowski sum $(1 - \lambda)\sigma + \lambda\tau$ of subsets $\sigma \subseteq P$ and $\tau \subseteq Q$. A *mixed subdivision* of $(1 - \lambda)P + \lambda Q$ is a dissection of $(1 - \lambda)P + \lambda Q$ into finitely many mixed cells such that any two intersect in common faces (possibly empty).

A two-dimensional mixed cell is *fine* if it is the Minkowski sum of either two edges or of a point and a triangle. In the first case, the cell is a parallelogram, in the second case the cell is a triangle. A mixed subdivision is *fine* if it contains only fine mixed cells.

4.2. Cayley embeddings. Let P and Q as above. Then the *Cayley embedding of P and Q* is the point configuration

$$\mathcal{C}(P, Q) := \{(p, 0) : p \in P\} \cup \{(q, 1) : q \in Q\} \subset \mathbb{R}^3.$$

For example, $P_n(\alpha)$ from above is a Cayley embedding for all α .

4.3. The Cayley trick. The Cayley trick states that for all P and Q as above, triangulations of $\mathcal{C}(P, Q)$ are in one-to-one correspondence with fine mixed subdivisions of $(1 - \lambda)P + \lambda Q$ for all $\lambda \in (0, 1)$. We will only need the fact that every triangulation of $\mathcal{C}(P, Q)$ induces a fine mixed subdivision of $(1 - \lambda)P + \lambda Q$ for all $\lambda \in (0, 1)$.

The correspondence is given by intersecting $\mathcal{C}(P, Q)$ by a horizontal hyperplane H_λ at height λ . The intersection of any tetrahedron in a triangulation of $\mathcal{C}(P, Q)$ with H_λ is a fine mixed cell in $((1 - \lambda)P + \lambda Q) \times \{\lambda\} \subset \mathbb{R}^3$. Since intersection with affine hyperplanes preserves face relations, the set of all fine mixed cells so obtained yields a fine mixed subdivision of $(1 - \lambda)P + \lambda Q$.

Applied to $P_n(\alpha)$ this means: each triangulation of $P_n(\alpha)$ induces a fine mixed subdivision of $S_n(\alpha, \lambda) := (1 - \lambda)C_n + \lambda C_n(\alpha)$ for every $\lambda \in (0, 1)$. In summary, we have the following correspondences between objects in the Cayley embedding and the Minkowski sum:

$P_n(\alpha)$	$S_n(\alpha, \lambda)$
tetrahedra	fine mixed polygons
tetrahedra with a triangle in the bottom or the top	fine mixed triangles
tetrahedra with edges both in the top and in the bottom	fine mixed parallelograms
non-horizontal triangles	fine mixed edges
non-horizontal edges	fine mixed points
orientation	counter-clockwise orientation

Since the Minkowski sum lives in one dimension less than the Cayley embedding, we rather work with $S_n(\alpha, \lambda)$.

5. THE PROOFS

Let $\alpha \geq 0$ be small enough such that $P_n(\alpha)$ is a prism or a twisted prism. Fix $\varepsilon \in (0, 1)$ such that $\|\varepsilon(v_j - v_i)\| < \|(1 - \varepsilon)(w_j - w_i)\|$ for all $i, j = 0, 1, \dots, n-1$. (All following considerations are also true for arbitrary $\varepsilon \in (0, 1)$; the choice of a small ε makes some arguments more transparent, though.) In particular, $S_n(\alpha) := S_n(\alpha, 1 - \varepsilon) = \varepsilon P_n + (1 - \varepsilon)P_n(\alpha)$ does not contain multiple points. For brevity, we will use the notation (i, j) for the point $\varepsilon v_i + (1 - \varepsilon)w_j$, $i, j = 0, 1, \dots, n-1$.

5.1. Some notions and notation. Consider mixed edges. All mixed edges are, by definition, Minkowski sums of either a point and an edge or of an edge and a point. In our notation, they are of the form $(e, i) := \{(k, i), (l, i)\}$ or of the form $(j, e) := \{(j, k), (j, l)\}$ for some edge $e = \{k, l\}$ in C_n or $C_n(\alpha)$, resp.

The following notions are motivated by regarding ε as being small. We highlight the most important one as a definition.

Definition 1 (Short and Long Edges). Call a mixed edge *short* if it is of the form (e, i) , call it *long* otherwise. The short mixed edge $e_i := \{(i, i), (i+1, i)\}$ is called *special*. \square

The special edges are interesting in S_n because – via the Cayley trick – they correspond to triangles that are incompatible with the cyclic set of diagonals D_c in P_n . Moreover, they are interesting in $S_n(\alpha)$ for $\alpha > 0$ because the cyclic set of tetrahedra T_c covers the corresponding triangles in $P_n(\alpha)$ so that in any triangulation containing T_c no other cell can use them.

For $i = 0, 1, \dots, n-1$, there are the convex sub- n -gons $(C_n, i) := \varepsilon C_n + (1 - \varepsilon)w_i$ and $(i, C_n(\alpha)) := \varepsilon v_i + (1 - \varepsilon)C_n(\alpha)$ in S_n . By construction, all (C_n, i) are translates of C_n , and all $(i, C_n(\alpha))$ are translates of $C_n(\alpha)$, which itself is an angle-preserving image of C_n under a (small) rotation that we call $r(\alpha)$. The *long* translation that shifts (C_n, i) to (C_n, j) along the long edge $\{(i, i), (i, j)\}$ is denoted by T_{ij} ; the *short* translation that moves $(i, C_n(\alpha))$ to $(j, C_n(\alpha))$ along the short edge $\{(i, i), (j, i)\}$ is denoted by t_{ij} .

Call the n -gons (C_n, i) *small* and the n -gons (j, C_n) *large*. Similarly, we call a mixed triangle with only short edges *small*; we call a mixed triangle with only long edges *large*. By definition of the Minkowski sum, each mixed triangle is either small or large. We can regard short mixed edges as edges that have both end points in the same small sub- n -gon. Since C_n is convex, no line spanned by a short edge e in (C_n, i) cuts the interior of any other short edge in (C_n, i) . The special short mixed edge e_i lies in the boundary of $S_n(\alpha)$. Figure 1 illustrates the setup.




FIGURE 1. Cutting the Cayley embedding of two n -gons with a horizontal hyperplane close to the top yields their Minkowski sum scaled as in $S_n(\alpha)$; the cyclic set of diagonals and the special edges are drawn thicker

5.2. Roadmap of the proofs. Note that any triangulation of P_n that uses the cyclic set of diagonals induces a mixed subdivision M of S_n in which no special edge e_i is used. Consider any non-special short edge e in M in some small n -gon (C_n, i) . Then the “region” between e and e_i must be covered by M somehow. We want to show that this cannot be accomplished unless at least one special edge is used. We even show that at least one special edge must be used as an edge of some mixed triangle (Theorem 3).

How can the region between e and e_i be subdivided? There must be a cell adjacent to e on the same side as e_i . If we use a mixed triangle, i.e., a small triangle, then we harvest new short edges in the same small n -gon. One of these new short edges is “closer” to e_i in a sense to be defined precisely below, and we can proceed. If we use a mixed

parallelogram then there is another short edge e' opposite to e in some other small n -gon (C_n, j) at a “partner vertex” j of e . But the “regions” containing potential partner vertices for e' towards e_j will turn out to be strictly smaller than for e .

But what happens if we use a mixture of mixed triangles and parallelograms? In fact, both ideas from above can be merged by using a certain lexicographic partial order on short edges, in which the short edges that are hit by “chasing the mixed subdivision M towards special short edges” are strictly decreasing. This shows that not all special short edges can be avoided by M .

We can make this idea precise for both the prism and the twisted prism. In the latter case, it is no surprise that even all special edges must be used, since they are boundary edges of $S_n(\alpha)$. However, using the cyclic set of tetrahedra means covering all special short edges by parallelograms, and we will show that at least one of them must be in a small triangle.

In the sequel, we will formalize these arguments in order to obtain rigorous proofs of Theorems 1 and 2.

5.3. Ordering short mixed edges. For the following, let e be a short edge in (C_n, i) . We want to give an orientation to the halfplanes separated by the line $l(e)$ spanned by e . If $e = e_i$, then we make use of the fact that e_i is in the boundary of S_n , thus $l(e)$ is a supporting hyperplane for S_n . Therefore, we can define the positive side $l(e)^+$ of e to be the halfplane not containing S_n . If $e \neq e_i$, we define the positive side $l(e)^+$ of e to be the halfplane containing e_i . This idea of investigating the subdivision between e and e_i can now be formulated as looking at cells on the positive side of $l(e)$.

The following is a simple observation.

Lemma 1. *Let σ be a mixed parallelogram in $S_n(\alpha)$ with short edges e and e' . Then:*

- (i) *If σ is on the positive sides or on the negative sides of both of its short edges then $l(e)$ and $l(e')$ have opposite orientations.*
- (ii) *If σ is on the positive side of e and on the negative side of e' , or vice versa, then $l(e)$ and $l(e')$ have parallel orientations. \square*

One of the cases mentioned in Lemma 1 can actually never occur. This will allow us to keep on finding new cells on the positive sides of short edges.

Lemma 2 (Orientation Lemma). *There is no fine mixed 2-cell σ in S_n on the positive side of all of its short edges.*

Remark 2. The correctness of the Orientation Lemma heavily depends on the congruence of the top/bottom polygons of $P_n(\alpha)$ and on the restriction of α . That the lemma is false in even slightly more general situations can be seen in the example in Figure 2.

Proof. Assume, for the sake of contradiction, that σ is a mixed 2-cell in S_n lying on the positive side of all of its short edges. Since σ contains the short edge e , it must be either a small triangle or a parallelogram.

Consider the case where σ is a small triangle on the positive side of all of its edges. The special edge e_i cannot be an edge of σ , since σ is contained in $\text{conv} S_n$, and $l(e_i)^+$ was defined to be the side of $l(e_i)$ not containing S_n . By definition of the orientations of short




FIGURE 2. Parallelograms which are on the positive sides of both of their short edges exist when α is too large; in the picture $\alpha = \frac{\pi}{3}$. However, it can be seen that the bad parallelogram flips its orientation when $P_4(\alpha)$ is un-twisted

edges other than e_i , we conclude that e_i must be contained in σ . Since (C_n, i) is convex, this can only be the case if e_i is an edge of σ : contradiction.

Therefore, σ must be a parallelogram lying on the positive sides of both of its short edges e in (C_n, i) and e' in (C_n, j) for some $i, j \in \{0, 1, \dots, n-1\}$. We first consider this in the case of the prism, i.e., when $\alpha = 0$. We will also include the degenerate case, i.e., where σ is a line segment, into our considerations. Since $\sigma \subset l(e)^+ \cap l(e')^+$, the orientations of e and e' must be opposite (Lemma 1). In terms of translations, $T_{ij}(l(e)^+) = l(e')^-$ and $T_{ji}(l(e')^+) = l(e)^-$. By definition of the orientation, e_i is on the positive side of e , and hence $(i, i) \in l(e)^+$. Similarly, $(j, j) \in l(e')^+$. This implies

$$(i, i) \in l(e)^+, \quad (1)$$

$$(j, i) = T_{ji}(j, j) \in T_{ji}(l(e')^+) = l(e)^-, \quad (2)$$

$$(i, j) = T_{ij}(i, i) \in T_{ij}(l(e)^+) = l(e')^-, \quad (3)$$

$$(j, j) \in l(e')^+. \quad (4)$$

These are necessary conditions for a non-degenerate σ being on the positive side of both of its short edges. While being on the positive side of short edges does not make sense for

degenerate σ , Conditions (1) through (4) have a meaning in the degenerate case as well. For further reference, we call these necessary conditions the *orientation conditions*.

Since $\alpha = 0$, the points (i, i) , (j, i) , (i, j) , and (j, j) lie on a straight line ℓ . Since ε is very small, the points appear on ℓ in the order (i, i) , (j, i) , (i, j) , and (j, j) . This tells us that ℓ starts in $l(e)^+$, enters $l(e)^-$, and then returns into $l(e)^+$. This implies that $\ell = l(e)$. By the symmetric argument, also $\ell = l(e')$. Therefore, σ is a segment. Moreover, its short edges are actually $e = \{(i, i), (j, i)\}$ and $e' = \{(i, j), (j, j)\}$ because the points in (C_n, i) are in strictly convex position.

This shows that a non-degenerate σ cannot exist in the prism. Moreover, we have learned the following useful fact: if the points (i, i) , (j, i) , (i, j) , and (j, j) satisfy the orientation conditions (1) through (4) for the short edges e and e' of some (possibly degenerate) parallelogram σ then $\sigma = \{(i, i), (j, i), (i, j), (j, j)\}$.

Since σ cannot exist in the prism, consider the case where $\alpha > 0$ so that $P_n(\alpha)$ is still a twisted prism. That means, no full-dimensional tetrahedron in P_n switches orientation during the twisting towards $P_n(\alpha)$. That implies that no full-dimensional parallelogram in $S_n(0)$ changes its orientation w.r.t. its short edges (by the Cayley trick correspondence in Section 4.3; easy exercise in linear algebra).

Now, untwist $P_n(\alpha)$, and hence σ . Then, σ must degenerate to a segment in P_n . During the untwist, for all $\alpha > 0$ the points (i, i) , (j, i) , (i, j) , and (j, j) must always satisfy the orientation conditions. Since the conditions define a closed space and untwisting changes all data continuously in α , they must also hold in the degenerate position $\alpha = 0$. Hence, σ must be of the form $\{(i, i), (j, i), (i, j), (j, j)\}$ for some $i, j \in \{0, 1, \dots, n-1\}$. In particular, $e = \{(i, i), (j, i)\}$.

We finally show that during the twist, σ folds up in the “wrong” direction. Consider the order of the short edges incident to (i, i) counter-clockwise starting at an edge of S_n . In this order e_i is the first edge, by definition. Twisting P_n again counter-clockwise by α will turn the slope of the short edge $e = \{(i, i), (j, i)\}$ counter-clockwise into the slope of the long edge $\{(i, i), (i, j)\}$. Therefore, the long edge $\{(i, i), (i, j)\}$ and the special short edge e_i are on different sides of e . This means, σ lies on the negative side of e : contradiction. \square




FIGURE 3. Primary index $\text{ind}_1(e)$ of a short edge e

The following quantity defines how close a short edge is to the corresponding special short edge. See Figure 3 for an illustration.

Definition 2 (Primary Index). We define the *primary index* $\text{ind}_1(e)$ of any short edge e by

$$\text{ind}_1(e) := \text{vol}(\text{conv}(C_n, i) \cap l(e)^+) \quad \square$$

We now turn our attention to measuring how many short partner edges a short edge can find to build a parallelogram on its positive side. Consider the unique line $l(e, i)$ parallel to e through (i, i) . Let $l(e, i, \alpha)$ be the line that is obtained from $l(e, i)$ by a rotation by $-\alpha$ around (i, i) . Its orientation is obtained by rotating the orientation of $l(e)$ by $-\alpha$ as well. The resulting positive halfplane defined by $l(e, i, \alpha)$ is called $l(e, i, \alpha)^+$.

Lemma 3 (Partner Lemma). *Let σ be a mixed parallelogram with short edges e and e' so that σ lies on the positive side of e . Assume, e lies in the small polygon (C_n, i) and e' lies in the small polygon (C_n, j) . Then (j, i) lies in the interior of $l(e, i, \alpha)^+$.*

Proof. Assume, for the sake of contradiction, that (j, i) lies in $l(e, i, \alpha)^-$. By definition, e_i is inside $l(e)^+$. Since e_i is a boundary edge of $S_n(\alpha)$, one of the long edges E of σ must separate e_i from σ . Let $(k, i) := E \cap e$, where $k = i$ is possible.

Let β be the angle from e to E around (k, i) . This angle is the same as the angle from $l(e, i)$ to $\{(i, i), (i, j)\}$ around (i, i) : the short translation t_{ki} moves (k, i) to (i, i) , E onto $\{(i, i), (i, j)\}$, and e into $l(e, i) \cap \text{conv } S_n(\alpha)$. There are two cases: either $0 < \beta < \pi$ or $-\pi < \beta < 0$.

If $0 < \beta < \pi$ then the slope of e turns counter-clockwise around (k, l) into the slope of E . Since σ , and hence E , are in $l(e)^+$, the interior of the positive side $l(e)^+$ of $l(e)$ can be characterized as follows: a point $x \in \mathbb{R}^2$ is in the interior of $l(e)^+$ if and only if the angle from e to $\{(k, i), x\}$ around (k, i) is in the interval $(0, \pi)$. Since the orientation of $l(e, i)$ is parallel to this, the analogous characterization holds for the interior of $l(e, i)^+$. The characterization of the interior of the positive side $l(e, i, \alpha)^+$ of $l(e, i, \alpha)$ is analogous.





FIGURE 4. The case $0 < \beta < \pi$ in the proof of the Partner Lemma

Let γ be the angle from $\{(i, i), (j, i)\}$ to $l(e, i, \alpha)$ around (i, i) . The assumption that (j, i) lies in $l(e, i, \alpha)^-$ can now be expressed as $-\gamma \in [-\pi, 0] \iff \gamma \in [0, \pi]$. The angle from

FIGURE 6. Secondary index $\text{ind}_2(e)$ of a short edge e

order that will always decrease when we “chase M along short edges towards special short edges”.

Definition 4. Let e and e' be short edges in M' . Then

$$e \prec e' : \iff \begin{cases} \text{either} & \text{ind}_1(e) < \text{ind}_1(e') \\ \text{or} & \text{ind}_1(e) = \text{ind}_1(e') \text{ and } \text{ind}_2(e) < \text{ind}_2(e'). \end{cases} \quad \square$$

The following lemma is the formalization of “chasing the mixed subdivision towards special short edges”.

Lemma 4 (Order Lemma). *Let e be a short edge in a mixed subdivision M of $S_n(\alpha)$. Then the following hold:*

- (i) $\text{ind}_1(e) \geq 0$ and $\text{ind}_2(e) \geq 0$.
- (ii) $\text{ind}_1(e) = 0$ if and only if $e = e_i$ for some $i = 0, 1, \dots, n-1$.
- (iii) If $e \neq e_i$ for all $i = 0, 1, \dots, n-1$, then there exists another short edge e' in M with $e' \prec e$; moreover, there exists a 2-cell σ such that both e and e' are short edges of σ , and σ is on the positive side of e and on the negative side of e' .

Proof. Assertions (i) and (ii) are by definition.

In order to prove (iii), consider a short edge e in M . Assume that e is in (C_n, i) and that $e \neq e_i$. Then the mixed subdivision M must contain cells that subdivide the convex hull of e and e_i . In particular, there must be a cell σ on the positive side of e . There are two cases: Either σ is a simplex containing only short edges inside (C_n, i) , or σ is a parallelogram containing two short and two long edges.

Case 1: The cell σ is a simplex with short edges. By construction, $l(e)^+$ contains σ . By Lemma 2, σ lies on the negative side of one of its short edges, say e' . Then $l(e')^+$ does not contain σ . Moreover, since (C_n, i) is convex, $l(e)$ and $l(e')$ do not cross inside $\text{conv}(C_n, i)$. Thus, $l(e')^+ \cap \text{conv}(C_n, i) \subseteq l(e)^+ \cap \text{conv}(C_n, i) \setminus \sigma$. Therefore, $\text{ind}_1(e') \leq \text{ind}_1(e) - \text{vol}(\sigma) < \text{ind}_1(e)$, whence $e' \prec e$.

Case 2: The cell σ is a parallelogram containing two short and two long edges. Consider the short edge e' in σ opposite to e . It lies in (C_n, j) for some $j = 0, 1, \dots, n-1$ with $j \neq i$.

We first prove that e and e' have the same primary index. By Lemma 2, σ lies on the negative side of e' . By construction, σ lies on the positive side of e . Therefore, by Lemma 1,

the parallel lines $l(e)$ and $l(e')$ have parallel orientations. That means, $T_{ij}(l(e)^+) = l(e')^+$. Because $T_{ij}(\text{conv}(C_n, i)) = \text{conv}(C_n, j)$, we conclude $\text{ind}_1(e') = \text{ind}_1(e)$.

Next, we show that the secondary index of e' is strictly smaller than that of e . By Lemma 3, (j, i) lies in the interior of $l(e, i, \alpha)^+$. This implies that $(j, j) = T_{ij}(j, i)$ lies in the interior of $T_{ij}(l(e, i, \alpha)^+)$. Since the parallel lines $l(e)$ and $l(e')$ have parallel orientations, the parallel lines $l(e, i, \alpha)$ and $l(e', j, \alpha)$ also have parallel orientations. Thus, $l(e', j, \alpha)^+$ is strictly contained in $T_{ij}(l(e, i, \alpha)^+)$. Therefore,

$$\begin{aligned} \text{ind}_2(e') &= \text{vol}(\text{conv}(C_n, j) \cap l(e', j, \alpha)^+) \\ &= \text{vol}(\text{conv } T_{ij}(C_n, i) \cap l(e', j, \alpha)^+) \\ &< \text{vol}(\text{conv } T_{ij}(C_n, i) \cap T_{ij}(l(e, i, \alpha)^+)) \\ &= \text{vol}(\text{conv}(C_n, i) \cap (l(e, i, \alpha)^+)) \\ &= \text{ind}_2(e'). \end{aligned}$$

This proves that $e' \prec e$, and (iii) is proven as well. \square

5.4. The neighborhood of special short edges. We are now in a position to prove the main property of mixed subdivisions of $S_n(\alpha)$.

Theorem 3. *Let $\alpha \geq 0$ such that $P_n(\alpha)$ is a prism or a twisted prism. Then every mixed subdivision M of $S_n(\alpha)$ contains at least one triangle one of whose edges is some special short edge.*

Remark 3. If α is too large then not only the Order Lemma is false but also Theorem 3, which can be seen in Figure 7. Theorem 2, however, might still be true for large α because the cyclic set of tetrahedra defines parallelograms that are incompatible with the parallelogram that is on the positive sides of both of its short edges in Figure 7. One could consider all $\alpha \geq 0$ for which the face lattice of $P_n(\alpha)$ equals the one of the twisted prism in our sense. Since the existence of triangulations depends on the orientations of tetrahedra (the oriented matroid) rather than on the face lattice, we decided not to investigate this any further. If the top and the bottom n -gons are not congruent, Theorem 3 – and even Theorem 1 – do not hold either, as can be seen in Figure 8.

Proof. Since every triangulation of $P_n(\alpha)$ induces a triangulation of its top and its bottom polygon, at least one short triangle must be used. Not all of its short edges can be edges of $S_n(\alpha)$. Therefore, there is a short edge having cells on both of its sides. Hence, there is at least on 2-cell that is on the positive side of some short edge. By Lemma 2, every such cell lies on the negative side of one of its other short edges.

Let σ be a cell on the positive side of its short edge e and on the negative side of its short edge e' such that e' is minimal w.r.t. “ \prec ”. Then, by Lemma 2(iii), e' is a special edge.

Every parallelogram σ with a special short edge e_i must lie on the negative side of e_i , since the positive side of e_i is outside $S_n(\alpha)$. Therefore, the parallelogram σ lies on the same side of e_i as (C_n, i) . Assume the opposite edge e of σ lies in (C_n, j) for some $j \in \{0, 1, \dots, n-1\}$. Then, by Lemma 1, σ lies on the opposite side of e as (C_n, j) . In




FIGURE 7. When α is too large (in the picture $\alpha = \frac{\pi}{3}$) then a mixed subdivision exists where no special edge is covered by a mixed triangle; the parallelogram of Figure 2 serves as kind of an adapter between two part of the subdivision that would be incompatible otherwise. This mixed subdivision disappears when $P_4(\alpha)$ is untwisted. The indicated mixed subdivision does, however, not contradict the statement in Theorem 2 for larger α , since it does not use the parallelograms corresponding to the cyclic set of tetrahedra

particular, σ lies on the opposite side of e as e_j , which means, σ lies also on the negative side of e . \square

5.5. The proof of Theorem 1 (prism). For the sake of contradiction, assume that there is a triangulation T of P_n that uses the cyclic set D_c of diagonals. Using the Cayley trick, T induces a fine mixed subdivision M of S_n that uses, among others, the set of points $(i, i+1)$ for all $i = 0, 1, \dots, n-1$, corresponding to the cyclic set of diagonals (labels again regarded modulo n). The triangles in the quadrilateral facets of P_n induce the mixed edges $\{(i, i), (i, i+1)\}$ in the boundary of S_n . They already cover the whole boundary of S_n . Thus, the special edges $e_i := \{(i, i), (i+1, i)\}$ in the boundary of S_n , which correspond to the reverse cyclic set of diagonals in the quadrilateral facets of P_n , are not used in M . However, by Theorem 3, at least one e_i must be in M : contradiction. \square




FIGURE 8. Congruence of top and bottom n -gons is important: even if the top and the bottom n -gon of a Cayley embedding of two n -gons are normally equivalent, there may be triangulations using the cyclic set of diagonals of the resulting combinatorial polygonal prism; the figure shows the corresponding mixed subdivision; note that indeed no special edge is used

5.6. The proof of Theorem 2 (twisted prism). For the sake of contradiction, assume that there is a triangulation T of P_n that uses the cyclic set T_c of tetrahedra. Construct the corresponding mixed subdivision M of $S_n(\alpha)$. The set M_c of mixed cells corresponding to T_c are parallelograms that cover all the special edges e_i . Therefore, there can be no other cell that contains a special edge. By Theorem 3, there must be at least one mixed triangle containing a special edge e_i : contradiction. \square

ACKNOWLEDGEMENTS

I thank Yuri Romanovsky for bringing up this problem and Günter Ziegler for communicating it to me.

REFERENCES

1. Frederick Bagemihl, *On indecomposable polyhedra*, American Mathematical Monthly **55** (1948), 411–413.
2. Anders Björner, Michel Las Vergnas, Bernd Sturmfels, Neil White, and Günter M. Ziegler, *Oriented matroids*, Encyclopedia of Mathematics, vol. 46, Cambridge University Press, Cambridge, 1993.
3. Birkett Huber, Jörg Rambau, and Francisco Santos, *Cayley embeddings, lifting subdivisions, and the Bohné-Dress theorem on zonotopal tilings*, Journal of the European Mathematical Society **2** (2000), 179–198.
4. Nels Johann Lennes, *Theorems on the simple finite polygon and polyhedron*, American Journal of Mathematics **33** (1911), 37–62.
5. Jim Ruppert and Raimund Seidel, *On the difficulty of triangulating three-dimensional nonconvex polyhedra*, Discrete & Computational Geometry **7** (1992), 227–253.
6. Erich Schönhardt, *Über die Zerlegung von Dreieckspolyedern in Tetraeder*, Mathematische Annalen **89** (1927), 309–312.

JÖRG RAMBAU, ZUSE-INSTITUTE BERLIN, TAKUSTR. 7, 14195 BERLIN, GERMANY
E-mail address: rambau@zib.de