

Combinatorics of Beacon-based Routing in Three Dimensions^{*}

Jonas Cleve^[0000-0001-8480-1726] (✉) and Wolfgang Mulzer^[0000-0002-1948-5840]

Institut für Informatik, Freie Universität Berlin, Berlin, Germany
{jonascleve,mulzer}@inf.fu-berlin.de

Abstract. A beacon is a point-like object which can be enabled to exert a magnetic pull on other point-like objects in space. Those objects then move towards the beacon in a greedy fashion until they are either stuck at an obstacle or reach the beacon's location. Beacons placed inside polyhedra can be used to route point-like objects from one location to another. A second use case is to cover a polyhedron such that every point-like object at an arbitrary location in the polyhedron can reach at least one of the beacons once the latter is activated.

The notion of beacon-based routing and guarding was introduced by Biro et al. [FWCG'11] in 2011 and covered in detail by Biro in his PhD thesis [SUNY-SB'13], which focuses on the two-dimensional case.

We extend Biro's result to three dimensions by considering beacon routing in polyhedra. We show that $\lfloor \frac{m+1}{3} \rfloor$ beacons are always sufficient and sometimes necessary to route between any pair of points in a given polyhedron P , where m is the number of tetrahedra in a tetrahedral decomposition of P . This is one of the first results that show that beacon routing is also possible in three dimensions.

1 Introduction

A *beacon* b is a point-like object in a polyhedron P which can be enabled to exert a magnetic pull on all points inside P . Those points then move in the direction in which the distance to b decreases most rapidly. As long as the distance decreases, points can also move along obstacles they hit on their way.

The resulting *attraction path* alternates between unrestricted movement inside P and restricted movement on the boundary of P . If the attraction path of a point p towards a beacon b ends in b we say that b *covers* p . On the other hand, p is *stuck* if it is in a position where it cannot decrease its distance to b .

A point p can be *routed via beacons* towards a point q if there exists a sequence of beacons $b_1, b_2, \dots, b_k = q$ such that b_1 covers p and b_{i+1} covers b_i for all $1 \leq i < k$. In our model at most one beacon can be enabled at any time and a point has to reach the beacon's location before the next beacon can be enabled.

The notion of beacon attraction was introduced by Biro et al. [4, 5] for two dimensions. This extends the classic notion of visibility [9]: the visibility region of a point is a subset of the attraction region of a point.

^{*} Supported in part by DFG grant MU 3501/1 and ERC StG 757609.

Here, we study the case of three-dimensional polyhedra. A three-dimensional polytope or *polyhedron* is a compact connected set bounded by a piecewise linear 2-manifold. The results in this work are based on the master’s thesis of the first author [8] in which various aspects of beacon-based routing and guarding were studied in three dimensions. Simultaneously, Aldana-Galván et al. [1,2] looked at orthogonal polyhedra and introduced the notion of *edge beacons*.

For two dimensions, Biro [4] provided bounds on the number of beacons for routing in a polygon. He also showed that it is NP-hard and APX-hard to find a minimum set of beacons for a given polygon such that it is possible to (a) route between any pair of points, (b) route one specific source point to any other point, (c) route any point to one specific target point, or (d) cover the polygon.

It is easy to reduce the two-dimensional problems to their three-dimensional counterparts by lifting the polygon into three dimensions. It thus follows that the corresponding problems in three dimensions are also NP-hard and APX-hard. More details can be found in [8, Chapter 4].

2 Preliminary Thoughts on Tetrahedral Decompositions

To show an upper bound on the number of beacons necessary to route between any pair of points in two dimensions Biro et al. [5] look at a triangulation of the polygon. They show that for every two additional triangles at most one beacon is needed. Even though there is a slight flaw in the case analysis of their proof, this can be easily repaired, see [8, Chapter 3.1] for more details and the working proof. We extend this approach to three dimensions by looking at the decomposition of a polyhedron into tetrahedra.

The tetrahedral decomposition is no general solution: Lennes [10] has shown in 1911 that a polyhedron cannot, in general, be decomposed into tetrahedra if no additional vertices are allowed. The problem of deciding whether such a decomposition exists is, in fact, NP-complete as shown by Ruppert and Seidel [11].

In the two-dimensional case, every simple polygon with n vertices has a triangulation with exactly $n - 2$ triangles; a polygon with h holes has a triangulation of $n - 2 + 2h$ triangles. In contrast, for three dimensions the number of tetrahedra in a tetrahedral decomposition is not directly related to the number of vertices. Chazelle [7] showed that for arbitrary n there exists a polyhedron with $\Theta(n)$ vertices for which at least $\Omega(n^2)$ convex parts are needed to decompose it. Naturally, this is also a worst-case lower bound on the number of tetrahedra. On the other hand, Bern and Eppstein [3, Theorem 13] show that any polyhedron can be triangulated with $\mathcal{O}(n^2)$ tetrahedra with the help of $\mathcal{O}(n^2)$ *Steiner points*. Furthermore, it is clear that every tetrahedral decomposition consists of at least $n - 3$ tetrahedra.

One polyhedron can have different tetrahedral decompositions with different numbers of tetrahedra. An example of such a polyhedron is a triangular bipyramid which can be decomposed into two or three tetrahedra, see [11, p. 228]. Due to this, we will prove bounds on the number of beacons needed for routing relative to the number of tetrahedra m rather than the number of vertices n . Since we accept

any kind of decomposition and do not have any general position assumption tetrahedral decompositions with Steiner points are allowed.

To successfully apply the ideas for two dimensions to three dimensions we need the following preliminary definition and lemma.

Definition 2.1 (Dual graph of tetrahedral decompositions). *Given a polyhedron with a tetrahedral decomposition $\Sigma = \{\sigma_1, \dots, \sigma_m\}$ into m tetrahedra, its dual graph is an undirected graph $D(\Sigma) = (V, E)$ where*

- (i) $V = \{\sigma_1, \dots, \sigma_m\}$ and
- (ii) $E = \{\{\sigma_i, \sigma_j\} \in \binom{V}{2} \mid \sigma_i \text{ and } \sigma_j \text{ share exactly one triangular facet}\}$.

Observation 2.2. Unlike in two dimensions, the dual graph of a tetrahedral decomposition is not necessarily a tree. We can still observe that each node in the dual graph has at most 4 neighbors—one for each facet of the tetrahedron.

Lemma 2.3. *Given a tetrahedral decomposition Σ of a polyhedron together with its dual graph $D(\Sigma)$ and a subset $S \subseteq \Sigma$ of tetrahedra from the decomposition whose induced subgraph $D(S)$ of $D(\Sigma)$ is connected, then*

- (i) $|S| = 2$ implies that the tetrahedra in S share one triangular facet,
- (ii) $|S| = 3$ implies that the tetrahedra in S share one edge, and
- (iii) $|S| = 4$ implies that the tetrahedra in S share at least one vertex.

Proof. We show this separately for every case.

- (i) This follows directly from Definition 2.1.
- (ii) In a connected graph of three nodes there is one node neighboring the other two. By Definition 2.1 the dual tetrahedron shares one facet with each of the other tetrahedra. In a tetrahedron every pair of facets shares one edge.
- (iii) By case (ii) there is a subset of three (connected) tetrahedra that shares one edge e . This edge is therefore part of each of the three tetrahedra. By Definition 2.1, the fourth tetrahedron shares a facet f with at least one of the other three (called σ). Since f contains three and e two vertices of σ they share at least one vertex. A depiction of the possible configurations of four tetrahedra can be seen in Fig. 1. □

3 An Upper Bound for Beacon-based Routing

After the preparatory work, we can now show an upper bound on the number of beacons needed to route within a polyhedron with a tetrahedral decomposition. The idea of the proof is based on the proof by Biro et al. [5] for (two-dimensional) polygons. We want to show the following

Hypothesis 3.1. Given a polyhedron P with a tetrahedral decomposition Σ with $m = |\Sigma|$ tetrahedra it is always sufficient to place $\lfloor \frac{m+1}{3} \rfloor$ beacons to route between any pair of points in P .

Since the proof is quite long and consists of many cases it is split up into various lemmas which are finally combined in Theorem 3.7.




Fig. 1. A polyhedron with a tetrahedral decomposition of four tetrahedra is in one of those three configurations. The shared vertex or edge is marked.

3.1 Preparation

Given the polyhedron P and a tetrahedral decomposition Σ with $m = |\Sigma|$ tetrahedra, we look at the dual graph $D(\Sigma)$ of the tetrahedral decomposition. For the rest of the section we want the dual graph to be a tree. This is possible by looking at a spanning tree T of $D(\Sigma)$ rooted at some arbitrary leaf node.

In the following, we will place beacons depending only on the neighborhood relation between tetrahedra. If T is missing some edge $\{u, v\}$ from $D(\Sigma)$ we “forget” that tetrahedra u and v are neighbors, i.e., share a common facet. We have less information about a tetrahedron’s neighborhood and thus we might place more beacons than needed—but never less.

Note 3.2. In the following we will refer to nodes of T as well as their corresponding tetrahedra with σ_i . It should be clear from the context when the node and when the tetrahedron is meant—if not, it is indicated.

The main idea of the proof is as follows: In a recursive way we are going to place a beacon and remove tetrahedra until no tetrahedra are left. As will be shown, for every beacon we can remove at least three tetrahedra which yields the claimed upper bound. We will show this by induction and start with the base case:

Lemma 3.3 (Base case). *Given a polyhedron P with a tetrahedral decomposition Σ with $m = |\Sigma| \leq 4$ tetrahedra it is always sufficient to place $\lfloor \frac{m+1}{3} \rfloor$ beacons to route between any pair of points in P .*

Proof. If $m = 1$ then P is a tetrahedron and due to convexity no beacon is needed.

If $2 \leq m \leq 4$ we can apply Lemma 2.3 which shows that all tetrahedra share at least one common vertex v . We place the only beacon we are allowed to place at v . Then v is contained in every tetrahedron and thus, by convexity, every point in P can attract and be attracted by a beacon at v . \square




Fig. 2. The possible configurations in the first part of the inductive step.

We can now proceed with the inductive step, that is, polyhedra with a tetrahedral decomposition of $m > 4$ tetrahedra. Our goal is to place k beacons which are contained in at least $3k + 1$ tetrahedra and can therefore mutually attract all points in those tetrahedra. Afterwards, we will remove at least $3k$ tetrahedra, leaving a polyhedron with a tetrahedral decomposition of strictly less than m tetrahedra, to which we can apply the induction hypothesis. We then need to show how to route between the smaller polyhedron and the removed tetrahedra.

To do this, we look at a deepest leaf σ_1 of the spanning tree T . If multiple leaves with the same depth exist we choose the one whose parent σ_2 has the largest number of children, breaking ties arbitrarily. In Fig. 2 we can see different cases how the part of T which contains σ_1 and σ_2 might look like. We first concentrate on Figs. 2a to 2e and show for them the first part of the inductive step. Note that in all five cases there needs to be at least one additional root node—either because we have strictly more than four tetrahedra or because the tree is required to be rooted at a leaf node. The second part of the inductive step, namely Fig. 2f will be dealt with in Lemma 3.6.

Lemma 3.4 (Inductive step I). *Given a polyhedron P with a tetrahedral decomposition Σ with $m = |\Sigma| > 4$ tetrahedra and a spanning tree T of its dual graph $D(\Sigma)$ rooted at some arbitrary leaf node. Let σ_1 be a deepest leaf of T with*

the maximum number of siblings and let σ_2 be its parent. Assume furthermore that any of the following conditions holds:

- (i) σ_2 has three children σ_1 , σ_3 , and σ_4 (see Fig. 2a),
- (ii) σ_2 has two children σ_1 and σ_3 and a parent σ_4 (see Fig. 2b),
- (iii) σ_2 has one child σ_1 and is the only child of its parent σ_3 whose parent is σ_4 (see Fig. 2c),
- (iv) σ_2 has one child σ_1 and its parent σ_3 has two or three children of which one, σ_4 , is a leaf (Fig. 2d), or
- (v) σ_2 has one child σ_1 and its parent σ_3 has three children each of which has a single leaf child (Fig. 2e).

Then we can place one beacon b at a vertex of σ_1 which is contained in at least four tetrahedra. We can then remove at least three tetrahedra containing b without violating the tree structure of T and while there is at least one tetrahedron left in T which contains b .

Proof. We show this individually for the conditions.

- (i)–(iv) In all those cases the induced subgraph of the nodes σ_1 , σ_2 , σ_3 , and σ_4 is connected. We can then see with Lemma 2.3(iii) that the four tetrahedra share at least one vertex at which b is placed.
After that we either remove σ_1 , σ_3 , and σ_4 (case (i)); σ_1 , σ_2 , and σ_3 (cases (ii) and (iii)); or σ_1 , σ_2 , and σ_4 (case (iv)). In all of those cases only leaves or inner nodes with all their children are removed which means that the tree structure of T is preserved. Additionally, we only remove three of the four tetrahedra that contain b , thus, one of them remains in T .
- (v) Looking at Fig. 2e we see that we have three different sets, each containing σ_3 , a child σ_i of σ_3 , and σ_i 's child: $\{\sigma_1, \sigma_2, \sigma_3\}$, $\{\sigma_5, \sigma_4, \sigma_3\}$, and $\{\sigma_7, \sigma_6, \sigma_3\}$. When applying Lemma 2.3(ii), we see that each set shares one edge, giving us three edges of σ_3 . Since at most two edges in any tetrahedron can be disjoint, at least two of the given edges must share a common vertex. Without loss of generality let these be the edges shared by $\{\sigma_1, \sigma_2, \sigma_3\}$ and $\{\sigma_5, \sigma_4, \sigma_3\}$. We can then place b at the shared vertex and afterwards remove σ_1 , σ_2 , σ_4 , and σ_5 . The beacon b is also contained in σ_3 which remains in T . \square

3.2 Special Cases in the Inductive Step

Until now, we have ignored the configuration in Fig. 2f. The problem here is that to remove the tetrahedra σ_1 to σ_5 we need to place two beacons. Placing two beacons but only removing five tetrahedra violates our assumption that we can always remove at least $3k$ tetrahedra by placing k beacons. If we removed σ_6 and σ_6 had additional children then T would no longer be connected which also leads to a non-provable situation. Thus, we need to look at the number and different configurations of the (additional) children of σ_6 .

Since there are many different configurations of σ_6 's children (and their subtrees) we decided to use a brute force approach to generate all cases we need




Fig. 3. One tetrahedron σ_6 with a subtree of five tetrahedra. Subfigures (b) and (c) depict configurations that satisfy cases (i) and (ii) of Lemma 3.5, respectively.

to look at. Afterwards we removed all cases where Lemma 3.4 can be applied and all cases where only the order of the children differed. This leaves us with nine different cases where (obviously) the subtree from Fig. 2f is always present. Thus we seek more information from this specific configuration.

Lemma 3.5. *Given a tetrahedral decomposition of six tetrahedra with the dual graph as depicted in Fig. 3a. Then at least one of the following holds:*

- (i) σ_1 to σ_5 share a common vertex, or
- (ii) $\sigma_3, \sigma_4, \sigma_5,$ and σ_6 share a common vertex v ; $\sigma_1, \sigma_2, \sigma_3,$ and σ_6 share a common edge e ; and $v \cap e = \emptyset$.

Proof. We first define $S_1 = \{\sigma_3, \sigma_4, \sigma_5, \sigma_6\}$ and $S_2 = \{\sigma_1, \sigma_2, \sigma_3, \sigma_6\}$. We observe that by Lemma 2.3 each set shares at least a vertex, but can also share an edge. We distinguish the cases by the shared geometric object:

- If both S_1 and S_2 each share an edge, case (i) holds. Each such edge needs to be part of the triangular facet which connects σ_3 and σ_6 . Thus both edges share a common vertex.
- If just one of the two sets shares an edge e and the other shares only a vertex v there are two trivial cases: If $v \cap e = \emptyset$ then case (ii) is true—see Fig. 3c for an example. On the other hand, if $v \cap e = v$ then case (i) holds.
- The last case is the one in which each of the sets shares only a vertex. The situation can be seen in Fig. 3b. First look at the vertex v of σ_3 not contained in the facet shared by σ_3 and σ_6 , i.e., $v \notin \sigma_3 \cap \sigma_6$. In the figure v is marked in orange. We observe that then all neighbors of σ_3 except σ_6 contain v . Thus, σ_2 contains v and three of its four facets are incident to v . One of the facets is the shared facet with σ_3 , but the other two are where σ_1 could be placed. σ_1 cannot be located at the fourth facet of σ_2 , since it would then share an edge with σ_3 and σ_6 which is covered in the previous cases.




Fig. 4. All “nontrivial” configurations of children of σ_6 . The tree in (a) is a subtree of all configurations. In all cases σ_6 has no other children than the ones shown here. Furthermore by the requirement that T is rooted at a leaf node, σ_6 needs to have an additional parent (except for case (a)).

Therefore, σ_1 contains v and with the same argument the same holds true for σ_5 . Then case (i) holds. \square

Lemma 3.6 (Inductive step II). *Given a polyhedron P with a tetrahedral decomposition Σ with $m = |\Sigma| > 4$ tetrahedra and a spanning tree T of its dual graph $D(\Sigma)$ rooted at some arbitrary leaf node. Let $T' \subseteq T$ be a subtree of T with height 3 for which Lemma 3.4 cannot be applied. (See Fig. 4 for all possible cases.)*

In T' we can then place $k \geq 2$ beacons for which it holds that the beacons are contained in at least $3k + 1$ tetrahedra and the graph of the beacons and the edges they are contained in is connected.

We can then remove at least $3k$ tetrahedra from T' , each of which contains a beacon, without violating the tree structure of T . After removal there is at least one tetrahedron left in T which contains one of the beacons.

Due to space constraints we omit the proof here. It involves analyzing each of the nine cases individually and carefully applying Lemma 3.5. The full proof can however be found in [8, Lemma 5.9, pp. 34–37].

Lemma 3.4 shows the inductive step for all subtrees with height at most 2 and Lemma 3.6 enumerates all combinatorically different subtrees with height exactly 3. Since in both lemmas the height of the subtree decreases by at least 1 we can always apply either of them.

3.3 Conclusion

We can now restate Hypothesis 3.1 as a theorem:

Theorem 3.7 (Upper bound). *Given a polyhedron P with a tetrahedral decomposition Σ with $m = |\Sigma|$ tetrahedra it is always sufficient to place $\lfloor \frac{m+1}{3} \rfloor$ beacons to route between any pair of points in P .*

Proof. We show this by induction. The base case is shown by Lemma 3.3. We assume that the induction hypothesis (Hypothesis 3.1) holds for all polyhedra with a tetrahedral decomposition with strictly less than m tetrahedra. We then show that it also holds for tetrahedral decompositions Σ with exactly m tetrahedra.

Look at a spanning tree T of the dual graph $D(\Sigma)$ of the tetrahedral decomposition Σ which is rooted at an arbitrary leaf node. Let σ_1 be a deepest leaf node and if σ_1 is not unique choose the one with the largest number of siblings, breaking ties arbitrarily. We can then apply either Lemma 3.4 or Lemma 3.6 and know at least the following:

- (i) We have placed $k \geq 1$ beacons and removed at least $3k$ tetrahedra.
- (ii) Every removed tetrahedron contains at least one beacon.
- (iii) The induced subgraph of the placed beacons on the vertices and edges of the polyhedron is connected.
- (iv) There is at least one beacon b contained in the remaining polyhedron P' .

From (i) it follows that the new polyhedron P' has a tetrahedral decomposition of $m' \leq m - 3k$ tetrahedra. We can then apply the induction hypothesis for P' . Thus we only need to place $k' = \lfloor \frac{m'+1}{3} \rfloor \leq \lfloor \frac{m-3k+1}{3} \rfloor = \lfloor \frac{m+1}{3} \rfloor - k$ beacons in P' to route between any pair of points in P' . Since $k' + k = \lfloor \frac{m+1}{3} \rfloor$ we never place more beacons than we are allowed.

From the induction hypothesis and (iv) we conclude that we are especially able to route from any point in P' to the beacon b and vice versa, since b is contained in P' . With (ii) we know that for every point p in the removed tetrahedra there is a beacon b' such that p attracts b' and b' attracts p . Finally, with (iii) we know that we can route between all beacons we have placed. This especially means that we can route from every beacon to the beacon b which is inside P' and vice versa.

This completes the inductive step and thus, by induction, we have proved the theorem. \square

Observation 3.8. Placing $\max(1, \lfloor \frac{m+1}{3} \rfloor)$ beacons is always sufficient to cover a polyhedron with a tetrahedral decomposition with m tetrahedra. We need at least one beacon to cover a polyhedron and placing them as in the previous proof is enough.

4 A Lower Bound for Beacon-based Routing

We now want to show a lower bound for the number of beacons needed to route within polyhedra with a tetrahedral decomposition. To do this we first show a different lower bound proof for two dimensions which can then be easily applied to three dimensions.

As shown by Biro et al. [6], $\lfloor \frac{n}{2} \rfloor - 1$ is not only an upper but also a lower bound for the necessary number of beacons in simple polygons. The idea for the following construction is similar to the one used by Shermer [12] for the lower bound for beacon-based routing in orthogonal polygons. We first show the construction of a class of spiral-shaped polygons for which we will then show that $\lfloor \frac{n}{2} \rfloor - 1$ beacons are needed for a specific pair of points.

Definition 4.1. *For every $c \in \mathbb{N}^{\geq 1}$ and some small $0 < \delta < 1$ a c -corner spiral polygon is a simple polygon with $n = 2c + 2$ vertices. These vertices, given in counterclockwise order, are called $s, q_1, q_2, \dots, q_c, t, r_c, r_{c-1}, \dots, r_1$ and their coordinates are given in polar notation as follows:*

- $s = (1; 0\pi), t = (\lfloor \frac{c+1}{3} \rfloor + 1; (c+1) \cdot \frac{2}{3}\pi),$
- $q_k = (\lfloor \frac{k}{3} \rfloor + 1 + \delta; k \cdot \frac{2}{3}\pi)$ for all $1 \leq k \leq c,$ and
- $r_k = (\lfloor \frac{k}{3} \rfloor + 1; k \cdot \frac{2}{3}\pi)$ for all $1 \leq k \leq c.$

The two vertices r_k and q_k form the k -th corner. The trapezoids $\Delta r_k q_k q_{k+1} r_{k+1}$ for all $1 \leq k < c$ and the two triangles $\Delta s r_1 q_1$ and $\Delta t r_c q_c$ are each called a hallway.

An example for $c = 5$ can be seen in Fig. 5. There are five corners and we have already placed five beacons to be able to route from s to t .

Lemma 4.2 (Two-dimensional lower bound). *Given a c -corner spiral polygon c beacons are necessary to route from s to t .*

Proof. To show that we need c beacons we introduce some additional notational conventions as depicted in Fig. 6a. The exterior angle at each (reflex) vertex r_k is called α_k . We split each hallway into two parts. To achieve this, three rays starting at the origin with the angles $\frac{1}{3}\pi, \pi,$ and $\frac{5}{3}\pi$ (the dotted rays in Fig. 5) are drawn. Every hallway is divided by exactly one of the three rays and the intersection points with the polygon's boundary are called a_k and b_k : a_k is the intersection of one of the rays with the edge $r_{k-1}r_k$ and b_k with the edge $q_{k-1}q_k$.

We observe that, due to the triangular shape, the angle α_k is always strictly less than 90° , for every $1 \leq k \leq c$.

We now divide our polygon into $c + 2$ subpolygons C_0 to C_{c+1} at each pair a_k and b_k . This subdivision results in two triangles C_0 and C_{c+1} which contain s and t , respectively, and c subpolygons C_1 to C_c which are called *complete corners* and which all have the same structure. We show that every such complete corner needs to contain at least one beacon to be able to route from s to t . We look at one complete corner $C_k, 1 \leq k \leq c,$ shown in Fig. 6b.




Fig. 5. A 5-corner spiral polygon with $\delta = 0.4$ for which five beacons (marked in red) are necessary to route from s to t .

We want to route from C_{k-1} to C_{k+1} . To route any point from C_{k-1} (which all lie to the right of $a_k b_k$) towards b_{k+1} there has to be a beacon such that the shortest line segment of this beacon to the line segment $r_k a_k$ ends in r_k . Otherwise, an attracted point will get stuck on $r_k a_k$ because the shortest path ends somewhere on $r_k a_k$ (excluding r_k itself). In Fig. 6b we see the case where b_{k+1} is a beacon. Here d is a dead point with respect to b_{k+1} and no point from C_{k-1} will travel further into C_k when attracted by b_{k+1} .

The marked region in Fig. 6b is the region in which every point can attract at least all points on the line segment $a_k b_k$. Additionally, every point in the region can be attracted by a beacon somewhere in the region of C_{k+1} to the left of the line trough $a_{k+1} r_k$, i.e., the part of C_{k+1} before turning at r_{k+1} .

On the other hand, there is no better option than b_{k+1} for a beacon position outside of C_k . All other points in C_{k+1} lie to the right of the line through $b_{k+1} d$ and hence their respective dead point on $r_k r_{k-1}$ lies further away from r_k .

We can see that if the length of the hallways (the distance between r_k and r_{k+1}) is sufficiently large compared to their width (the distance between a_k and b_k) it is never possible for a point outside of C_k to be inside the marked region. Therefore it is not possible to route from somewhere inside C_{k-1} to somewhere inside C_{k+1} without an additional beacon inside C_k . \square




Fig. 6. A more detailed look at the parts of the spiral polygon.

We now apply the idea of this proof to three dimensions starting with the definition of the spiral polyhedron.

Definition 4.3. For every $c \in \mathbb{N}^{\geq 1}$ and some small $0 < \delta < 1$ a c -corner spiral polyhedron is a polyhedron with $n = 3c + 2$ vertices. These vertices are s and t as well as q_k , r_k , and z_k for all $1 \leq k \leq c$. The coordinates of s , t , q_k , and r_k are the same as in Definition 4.1 with their z -coordinate set to 0. The z_k are positioned above r_k or more formally $z_k := r_k + \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ for all $1 \leq k \leq c$.

The edges and facets of the polyhedron are given by the tetrahedral decomposition:

- The start and end tetrahedra are formed by $\Delta r_1 q_1 z_1 s$ and $\Delta r_c q_c z_c t$.
- The hallway between two triangles $\Delta r_k q_k z_k$ and $\Delta r_{k+1} q_{k+1} z_{k+1}$ consists of the three tetrahedra $\Delta r_k q_k z_k r_{k+1}$, $\Delta r_{k+1} q_{k+1} z_{k+1} q_k$, and $\Delta q_k z_k r_{k+1} z_{k+1}$.

The three vertices r_k , q_k , and z_k form the k -th corner.

Observation 4.4. The smallest c -corner spiral polyhedron with $c = 1$ consists of exactly two tetrahedra. For greater c we add exactly $c - 1$ hallways, each consisting of three tetrahedra. This means that a c -corner spiral polyhedron has a tetrahedral decomposition with $m = 3c - 1$ tetrahedra. It follows from Definition 4.3 that the number of tetrahedra relative to the number of vertices is $m = 3 \cdot \frac{n-2}{3} - 1 = n - 3$.

Lemma 4.5 (Lower bound). Given a c -corner spiral polyhedron P , c beacons are necessary to route from s to t .

Proof. We project P onto the xy -plane which results in a c -corner spiral polygon P' due to the construction of the c -corner spiral polyhedron. To P' we can apply

Lemma 4.2 where we showed that c beacons (placed in an area around each of the c corners) are sometimes necessary to route in P' .

As opposed to the polygon, the movement in the polyhedron is not constrained to the xy -plane. Additionally, beacons can be placed at locations which do not lie in the xy -plane. We need to show that this does not change the situation in a way so that less than c beacons are necessary.

First, note that, due to the construction in Definition 4.3, every cross section of the polyhedron parallel to the xy -plane yields a c -corner spiral polygon with different widths δ . For every such cross section, Lemma 4.2 tells us that to route only in this cross section, c beacons are needed.

Additionally, the hallway's inner boundary $r_k z_k r_{k+1} z_{k+1}$ is perpendicular to the xy -plane. This means that the movement of all points p which are attracted by a beacon b can be split into a xy -movement and a z -movement because the z -coordinate is not important for any movement along the inner boundary. Since each hallway is convex there is no other movement of a point p attracted by a beacon b which is constrained by the polyhedron's boundary ∂P . We can then only look at the xy -movement which again yields a two-dimensional situation to which Lemma 4.2 can be applied. \square

5 A Sharp Bound for Beacon-based Routing

Theorem 5.1. *Given a polyhedron P for which a tetrahedral decomposition with m tetrahedra exists, it is always sufficient and sometimes necessary to place $\lfloor \frac{m+1}{3} \rfloor$ beacons to route between any pair of points in P .*

Proof. In Theorem 3.7 we have shown that $\lfloor \frac{m+1}{3} \rfloor$ is an upper bound.

For any given m we can construct a c -corner spiral polyhedron P_m with $c = \lfloor \frac{m+1}{3} \rfloor$ corners for which, by Lemma 4.5, c beacons are necessary. The number of tetrahedra in P_m is $m' = 3c - 1$ (see Observation 4.4) and this is also the smallest number of tetrahedra in any tetrahedral decomposition of P_m : If there was a tetrahedral decomposition with less tetrahedra then by Theorem 3.7 less than c beacons would be needed which contradicts Lemma 4.5.

If $m' < m$, i.e. due to the flooring function the c -corner spiral contains one or two tetrahedra less than m , we add the missing tetrahedra as if constructing a $c + 1$ -corner spiral. This does not lead to less beacons being needed. \square

6 Conclusion

We have shown that the problem of finding a minimal beacon set in a polyhedron P to route between all pairs of points or all points and a specific point is NP-hard and APX-hard. This holds also true for the problem of finding a minimal beacon set to cover a polyhedron P .

We have shown that, given a tetrahedral decomposition of a polyhedron P with m tetrahedra it is always sufficient to place $\lfloor \frac{m+1}{3} \rfloor$ beacons to route between

any pair of points in P . We then gave a class of polyhedra for which this upper bound is always necessary.

A lot of questions which have been answered by various authors in two dimensions remain open for the three-dimensional case. They include learning about the complexity of finding an optimal beacon set to route between a given pair of points. Additional open questions are about attraction regions (computing the set of all points attracted by a single beacon) and beacons kernels (all points at which a beacon can attract all points in the polyhedron).

Furthermore Cleve [8] has shown that not all polyhedra can be covered by beacons placed at the polyhedron's vertices and Aldana-Galván et al. [1,2] showed that this is even true for orthogonal polyhedra. Given a tetrahedral decomposition of a polyhedron it remains open whether it is possible to cover a polyhedron with less than $\max(1, \lfloor \frac{m+1}{3} \rfloor)$ beacons as seen in Observation 3.8. It seems challenging to further look at the beacon-coverage problem in general polyhedra.

Acknowledgments We thank the anonymous reviewers for their thorough reading of the paper and helpful suggestions.

References

1. Aldana Galván, I., Álvarez Rebolgar, J.L., Catana Salazar, J.C., Marín Nevárez, N., Solís Villarreal, E., Urrutia, J., Velarde, C.: Beacon coverage in orthogonal polyhedra. In: 29th Canadian Conference on Computational Geometry (CCCG 2017). pp. 166–171. Ottawa (Jul 2017)
2. Aldana-Galván, I., Álvarez-Rebolgar, J.L., Catana-Salazar, J.C., Marín-Nevárez, N., Solís-Villarreal, E., Urrutia, J., Velarde, C.: Covering orthotrees with guards and beacons. In: XVII Spanish Meeting on Computational Geometry (XVII ECG). Alicante (Jun 2017)
3. Bern, M., Eppstein, D.: Mesh generation and optimal triangulation. *Computing in Euclidean geometry* 4, 47–123 (1995)
4. Biro, M.: Beacon-Based Routing and Guarding. Ph.D. thesis, State University of New York at Stony Brook (2013)
5. Biro, M., Gao, J., Iwerks, J., Kostitsyna, I., Mitchell, J.S.B.: Beacon-based routing and coverage. In: 21st Fall Workshop on Computational Geometry (FWCG 2011) (2011)
6. Biro, M., Gao, J., Iwerks, J., Kostitsyna, I., Mitchell, J.S.B.: Combinatorics of beacon-based routing and coverage. In: Proceedings of the 25th Canadian Conference on Computational Geometry (CCCG 2013). vol. 1, p. 3 (2013)
7. Chazelle, B.: Convex partitions of polyhedra: A lower bound and worst-case optimal algorithm. *SIAM Journal on Computing* 13(3), 488–507 (1984)
8. Cleve, J.: Combinatorics of Beacon-Based Routing and Guarding in Three Dimensions. Master's thesis, Freie Universität Berlin, Berlin, Germany (Mar 2017)
9. Ghosh, S.K.: *Visibility Algorithms in the Plane*. Cambridge University Press, Cambridge (2007)
10. Lennes, N.J.: Theorems on the simple finite polygon and polyhedron. *American Journal of Mathematics* 33(1/4), 37 (1911)

11. Ruppert, J., Seidel, R.: On the difficulty of triangulating three-dimensional nonconvex polyhedra. *Discrete & Computational Geometry* 7(3), 227–253 (1992)
12. Shermer, T.C.: A combinatorial bound for beacon-based routing in orthogonal polygons. arXiv preprint arXiv:1507.03509 (2015)