Combinatorics of Beacon-based Routing in Three Dimensions*

Jonas Cleve[1] and Wolfgang Mulzer[1]

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

Abstract

A beacon is a point-like object that 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.

The notion of beacon-based routing was introduced by Biro et al. [FWCG'11] in 2011 for two dimensions and covered in detail by Biro in his PhD thesis [SUNY-SB'13].

We extend Biro's results to polyhedra in three dimensions. 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 that, when enabled, exerts a magnetic pull on points inside P. The points then move in the direction in which the distance to b decreases most rapidly, possibly moving along obstacles. If an attracted point p ends its movement at b, we say that b *covers* p. A point p can be *routed via beacons* towards a point q if there exists a sequence of beacons $b_1, b_2, \ldots, b_k = q$ such that b_1 covers p and b_{i+1} covers b_i for all $1 \leq i < k$. The target $q = b_k$ is an *implicit* beacon, i.e., we need only $k - 1$ additional beacons. We can route *between two points* p and q if p can be routed via beacons towards q and vice versa. In our model, at most one beacon is 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. It extends the classic notion of visibility [8]: the visibility region of a point p is a subset of the attraction region of p.

Here, we study the case of three-dimensional polyhedra. A three-dimensional *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 [7] 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, i.e., every point in P is covered by at least one beacon.

It is easy to reduce the two-dimensional problems to their three-dimensional counterparts by lifting the polygon into three dimensions. Thus the corresponding problems in three dimensions are also NP-hard and APX-hard. More details can be found in [7, Chapter 4].

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

34th European Workshop on Computational Geometry, Berlin, Germany, March 21–23, 2018.
This is an extended abstract of a presentation given at EuroCG'18. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

2 Preliminary Thoughts on Tetrahedral Decompositions

In two dimensions, Biro et al. [5] look at a triangulation of a polygon to show that for every two additional triangles at most one beacon is needed. This yields the upper bound of $\lfloor \frac{n}{2} \rfloor - 1$ beacons. Even though there is a slight flaw in the case analysis of their proof, this can be easily repaired, see [7, Chapter 3.1] for more details and the working proof. We extend their approach to three dimensions by looking at the decomposition of a polyhedron into tetrahedra. However, Lennes [9] has shown that polyhedra exist which cannot be decomposed into tetrahedra without additional vertices. To decide whether such a decomposition (without additional vertices) exists was proved to be NP-complete by Ruppert and Seidel [10].

In two dimensions, every triangulation of a polygon with n vertices and h holes has exactly $n - 2 + 2h$ triangles ($h = 0$ for simple polygons). In three dimensions, however, Chazelle [6] showed that there exist polyhedra with $\Theta(n)$ vertices for which $\Omega\left(n^2\right)$ convex parts are needed in every decomposition. Bern and Eppstein [3, p. 52] show that all polyhedra can be triangulated with $\mathcal{O}\left(n^2\right)$ tetrahedra with the help of $\mathcal{O}\left(n^2\right)$ additional *Steiner points*.

Additionally, one polyhedron can have different tetrahedral decompositions with different numbers of tetrahedra, see [10, p. 228]. Our results will therefore be relative to the number of tetrahedra m rather than the number of vertices n. We do not assume general position and thus decompositions with Steiner points are also allowed.

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

▶ **Definition 2.1.** 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
(1) $V = \{\sigma_1, \dots, \sigma_m\}$ and
(2) $E = \{\{\sigma_i, \sigma_j\} \in \binom{V}{2} \mid \sigma_i$ and σ_j share exactly one triangular facet$\}$.

▶ **Observation 2.2.** *Unlike in two dimensions, the dual graph of a tetrahedral decomposition is not necessarily a tree. Still, each node in the dual graph has at most* 4 *neighbors.*

▶ **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*
(1) $|S| = 2$ *implies that the tetrahedra in* S *share one triangular facet,*
(2) $|S| = 3$ *implies that the tetrahedra in* S *share one edge, and*
(3) $|S| = 4$ *implies that the tetrahedra in* S *share at least one vertex.*

Proof. We show this seperately for every case.
(1) This follows directly from Definition 2.1.
(2) 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.
(3) By case (2), 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. ◀

3 An Upper Bound for Beacon-based Routing

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

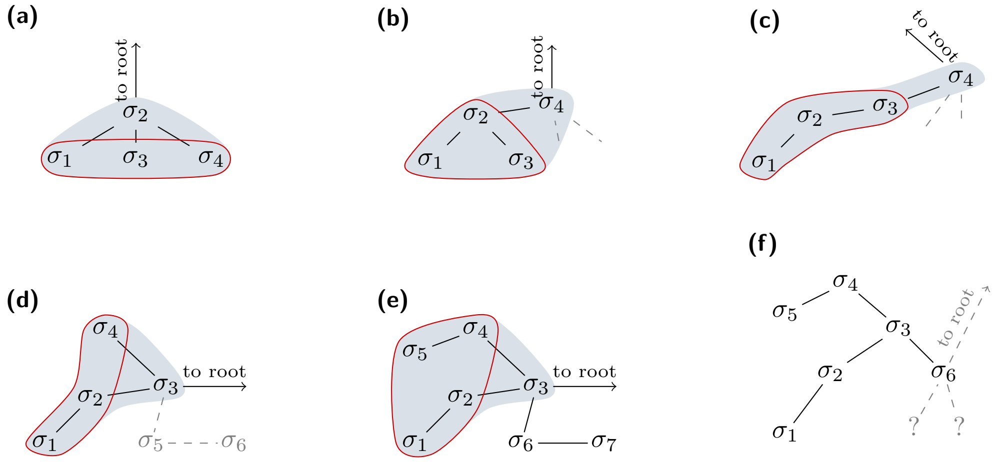

Figure 1 The possible configurations in the inductive step. The shaded region can be covered by one beacon, the circled tetrahedra are removed. Subfigure (f) has to be looked at independently.

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 always suffices to place $\left\lfloor \frac{m+1}{3} \right\rfloor$ beacons to route between any pair of points in P.*

Due to the length and number of cases, the proof is split up into various lemmas which are finally combined in Theorem 3.5.

Given the polyhedron P and a tetrahedral decomposition Σ with $m = |\Sigma|$ tetrahedra, we look at the dual graph $D(\Sigma)$ of the decomposition. We look at a spanning tree T of $D(\Sigma)$ rooted at some arbitrary leaf node because we want the dual graph to be a tree. This can only lead to more beacons being placed—never less. We will refer to nodes of T as well as their corresponding tetrahedra with σ_i—the meaning should be clear from the context.

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 each beacon we place, we can remove at least three tetrahedra. This yields the claimed upper bound. We will show this by induction on the number of tetrahedra. We first cover the base case:

▶ **Lemma 3.2.** *Given a polyhedron P with a tetrahedral decomposition Σ with $m = |\Sigma| \leq 4$ tetrahedra, it always suffices to place $\left\lfloor \frac{m+1}{3} \right\rfloor$ 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. Since v is contained in all tetrahedra and can thus attract and be attracted by all points in P, we are done by placing one beacon at v. ◀

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 that 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. 1, we can see different cases how the part of T that contains σ_1 and σ_2 might look like. We first show the inductive step for Figs. 1a to 1e. 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 inductive step for Fig. 1f will be dealt with in Lemma 3.4.

► **Lemma 3.3.** *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 further that either*

(1) the configuration in the neighborhood of σ_1 and σ_2 looks like any of Figs. 1a to 1d or

(2) the configuration in the neighborhood of σ_1 and σ_2 looks like Fig. 1e.

Then we can place one *beacon b at a vertex of σ_1 that 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 that contains b.*

Proof. We show this individually for the conditions.

(1) In all those cases the induced subgraph of the nodes σ_1, σ_2, σ_3, and σ_4 is connected. From Lemma 2.3(3) it follows that they share at least one vertex at which b is placed. Afterwards the circled tetrahedra are removed from T which preserves the tree structure of T. Additionally, at least one tetrahedron covered by b remains in T.

(2) Looking at Fig. 1e 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 $\{\sigma_7, \sigma_6, \sigma_3\}$.
By Lemma 2.3(2) 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 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. ◄

Until now, we have ignored the configuration in Fig. 1f. The problem here is that, in general, 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 to look at. Afterwards we removed all cases where Lemma 3.3 can be applied and all cases where only the order of the children differed. This leaves us with nine different cases and the following

► **Lemma 3.4.** *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 3.3 cannot be applied.*

In T' we can then place $k \geq 2$ beacons which are contained in at least $3k + 1$ tetrahedra and whose induced subgraph 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 which can, however, be found in [7, Lemma 5.9, pp. 34ff.]. There we also argue why either of the Lemmas 3.3 and 3.4 can always be applied.

We can now restate Hypothesis 3.1 as a theorem:

▶ **Theorem 3.5.** *Given a polyhedron P with a tetrahedral decomposition Σ with $m = |\Sigma|$ tetrahedra it always suffices 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 3.2.

Look at a spanning tree T of the dual graph $D(\Sigma)$ rooted at an arbitrary leaf node. Let σ_1 be a deepest leaf node with the largest number of siblings, breaking ties arbitrarily. We can then apply either Lemma 3.3 or Lemma 3.4 and know at least the following:

(1) We have placed $k \geq 1$ beacons and removed at least $3k$ tetrahedra.
(2) Every removed tetrahedron contains at least one beacon.
(3) The induced subgraph of the beacons on the vertices and edges of P, i.e. the graph which contains only the beacons as vertices and the edges between two beacons, is connected.
(4) There is at least one beacon b contained in the remaining polyhedron P'.

From (1) 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 $k' + k \leq \lfloor \frac{m+1}{3} \rfloor$ we never place more beacons than we are allowed.

With (2) to (4) we know that we can route from any point in P' to b, every point in the removed tetrahedra to at least one placed beacon, and between all placed beacons.

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

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. The idea for the following construction is similar to the one used by Shermer [11] for the lower bound for beacon-based routing in orthogonal polygons. We will present the results very briefly, for more details see [7].

We construct a class of outwards-spiraling polygons for which for every two additional vertices one additional beacon is needed. An example with $n = 12$ vertices and thus $\lfloor \frac{n}{2} \rfloor - 1 = 5$ needed beacons is shown in Fig. 2. It can easily be shown that it is not possible to route from s to t with less beacons, which gives the following

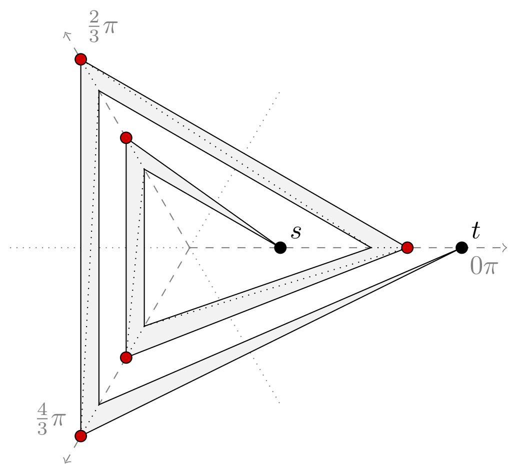

Figure 2 A spiral polygon with $n = 12$ vertices and $c = 5$ corners. Every such polygon needs $c = \frac{n}{2} - 1$ beacons to route from s to t.

▶ **Lemma 4.1.** *Given a spiral polygon with c corners and $n = 2c + 2$ vertices, c beacons are necessary to route from s to t.* ◀

The construction of the spiral polygon can be easily lifted to three dimensions by adding one vertex to each corner. This then directly leads to

▶ **Lemma 4.2.** *Given a spiral polyhedron with c corners and $n = 3c + 2$ vertices, c beacons are necessary to route from s to t.* ◀

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 $\left\lfloor \frac{m+1}{3} \right\rfloor$ beacons to route between any pair of points in P.*

Proof. In Theorem 3.5 we have shown that $\left\lfloor \frac{m+1}{3} \right\rfloor$ is an upper bound.

For any given m we can construct a spiral polyhedron P_m with $c = \left\lfloor \frac{m+1}{3} \right\rfloor$ corners for which, by Lemma 4.2, c beacons are necessary. The number of tetrahedra in P_m is $m' = 3c-1$ 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 3.5 less than c beacons would be needed which contradicts Lemma 4.2.

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 spiral polyhedron with $c + 1$ corners. This does not lead to less beacons being needed. ◀

References

1 I. Aldana Galván, J. L. Álvarez Rebollar, J. C. Catana Salazar, N. Marín Nevárez, E. Solís Villarreal, J. Urrutia, and C. Velarde. Beacon coverage in orthogonal polyhedra. In *29th Canadian Conference on Computational Geometry*, Ottawa, July 2017.

2 I. Aldana-Galván, J. L. Álvarez-Rebollar, J. C. Catana-Salazar, N. Marín-Nevárez, E. Solís-Villarreal, J. Urrutia, and C. Velarde. Covering orthotrees with guards and beacons. In *XVII Spanish Meeting on Computational Geometry (XVI ECG)*, Alicante, June 2017.

3 Marshall Bern and David Eppstein. Mesh generation and optimal triangulation. *Computing in Euclidean geometry*, 4:47–123, 1995.

4 Michael Biro. *Beacon-Based Routing and Guarding*. PhD thesis, State University of New York at Stony Brook, 2013.

5 Michael Biro, Jie Gao, Justin Iwerks, Irina Kostitsyna, and Joseph S. B. Mitchell. Beacon-based routing and coverage. In *21st Fall Workshop on Computational Geometry*, 2011.

6 Bernard Chazelle. Convex partitions of polyhedra: A lower bound and worst-case optimal algorithm. *SIAM Journal on Computing*, 13(3):488–507, 1984.

7 Jonas Cleve. Combinatorics of beacon-based routing and guarding in three dimensions. Master's thesis, Freie Universität Berlin, Berlin, Germany, March 2017.

8 Subir Kumar Ghosh. *Visibility Algorithms in the Plane*. Cambridge University Press, Cambridge, 2007. doi:10.1017/CBO9780511543340.

9 N. J. Lennes. Theorems on the simple finite polygon and polyhedron. *American Journal of Mathematics*, 33(1/4):37, 1911. doi:10.2307/2369986.

10 Jim Ruppert and Raimund Seidel. On the difficulty of triangulating three-dimensional nonconvex polyhedra. *Discrete & Computational Geometry*, 1992. doi:10.1007/BF02187840.

11 Thomas C. Shermer. A combinatorial bound for beacon-based routing in orthogonal polygons. *arXiv preprint arXiv:1507.03509*, 2015.