Cornell Theory Center / ACRI Technical Report CTC97TR266, 1/97

Quality Mesh Generation in Higher Dimensions*

Scott A. Mitchell† Stephen A. Vavasis‡

December 19, 1996

Abstract

We consider the problem of triangulating a d-dimensional region. Our mesh generation algorithm, called QMG, is a quadtree-based algorithm that can triangulate any polyhedral region including nonconvex regions with holes. Furthermore, our algorithm guarantees a bounded aspect ratio triangulation provided that the input domain itself has no sharp angles. Finally, our algorithm is guaranteed never to over-refine the domain in the sense that the number of simplices produced by QMG is bounded above by a factor times the number produced by any competing algorithm, where the factor depends on the aspect ratio bound satisfied by the competing algorithm. The QMG algorithm has been implemented in C++ and is used as a mesh generator for the finite element method.

1 Introduction

The *finite element method* refers to a family of numerical methods for solving boundary value problems and is used extensively in electromagnetics,

*This work supported in part by an NSF PYI award with matching funds from AT&T, Sun Microsystems, Tektronix and Xerox. Support was also received from Argonne National Laboratory.

†Computational Mechanics and Visualization Department, Sandia National Laboratories, Albuquerque, NM 87185, samitch@sandia.gov.

‡Department of Computer Science, Cornell University, Ithaca, NY 14853, vavasis@cs.cornell.edu. Part of this work was done while the author was visiting Argonne National Laboratory.

thermodynamics, structural analysis, acoustics, chemistry and astronomy. A crucial preprocessing step is mesh generation. A *mesh generator* is an algorithm for subdividing a finite subset of $\mathbb{R}^2$ or $\mathbb{R}^3$ into small convex cells, typically triangles or quadrilaterals in 2D and tetrahedra or hexahedra (brick shapes) in 3D.

We propose a mesh generation algorithm called QMG for nonconvex polyhedral regions in any dimension. QMG takes as input a representation of a polyhedral region in $\mathbb{R}^d$ and produces as output a simplicial complex that is a subdivision of the input region. QMG uses a quadtree technique: the domain is covered with a large d-dimensional cube, and then cubes are recursively split into 2^d subcubes until each subcube is triangulated.

For good accuracy bounds in the finite element method, it is necessary that the tetrahedra have bounded *aspect ratio.* The aspect ratio of a simplex is defined as its maximum side-length divided by its minimum altitude. For an analysis of the accuracy of the finite element method, see Johnson [9].

The mesh produced by QMG is guaranteed to have good aspect ratio. Let ρ_{QMG} be the worst aspect ratio among all simplices in the QMG triangulation of a particular input polyhedron P. Let $\rho_{\mathcal{S}}$ be the worst aspect ratio among all simplices in any other triangulation $\mathcal{S}$ of P, where $\mathcal{S}$ is produced by some other competing algorithm. Then Theorem 6 says that $\rho_{\mathrm{QMG}} \leq c\rho_{\mathcal{S}}$, where c is a universal constant, in the case $d = 2$ or $d = 3$. The technique used to prove this theorem is as follows. First, a lower bound is proved stating that any triangulation $\mathcal{S}$ of P must have at least one simplex with aspect ratio at least as large as $c/\theta(P)$, where $\theta(P)$ denotes the sharpest angle of P and c is some other constant. Then we prove that QMG's aspect ratio is bounded above by $c/\theta(P)$. In the case $d > 3$, a weaker version of this result is proved.

Our second main theorem is that the number of simplices generated by QMG is the smallest possible, i.e., the mesh is as coarse as possible, in the following sense. Let n_{QMG} be the number of simplices produced by QMG when applied to a particular polyhedral domain P, and let $n_{\mathcal{S}}$ be the number of simplices in some other triangulation $\mathcal{S}$ of P. Then $n_{\mathrm{QMG}} \leq f(d, \rho_{\mathcal{S}}) \cdot n_{\mathcal{S}}$, where f is some function of d, the dimension and of $\rho_{\mathcal{S}}$, the aspect ratio bound satisfied by the competing triangulation. In other words, n_{QMG} is much larger than $n_{\mathcal{S}}$ only in the case when $\mathcal{S}$ has simplices with poor aspect ratio. The precise values of the constants present in these two main results are not worked out explicitly in this paper but are expected to be quite large.

The importance of bounding the number of tetrahedra is as follows. The running time of the finite element method is a function of the number of nodes and elements in the triangulation. In particular, if n is the number of nodes (or elements—for bounded aspect ratio triangulations, the number of nodes and elements are within a constant factor of each other), then the running time of the finite element method is $O(n^\alpha)$, where α is at least 1 and depends on the method used for solving the sparse linear equations. Thus, there is a significant penalty for meshes with too many elements. On the other hand, small elements are necessary for high accuracy with the finite element method. Practitioners usually address this tradeoff by using meshes with varying degrees of refinement: such a mesh has small elements in the part of the domain of interest where high accuracy is desired, and larger elements are used elsewhere. Because QMG generates the coarsest mesh possible (up to the multiplicative factor $f(d, \rho_{\mathcal{S}})$), it can be used as the starting point for further refinement. Indeed, the implementation of QMG allows a user-specified refinement function to control the degree of refinement.

Our work is closely related to earlier work by Bern, Eppstein and Gilbert [4] who solved the corresponding problem for two-dimensional polygonal domains. These authors also used a quadtree approach, but the extension of their technique to higher dimensions is far from straightforward; the QMG algorithm differs in many ways from that earlier paper.

Other work on triangulation problems with optimality guarantees is the result of Baker, Grosse and Rafferty [1], whose algorithm triangulates 2D polygons with nonobtuse angles and Chew's [7] triangulation of 2D with guaranteed aspect ratio using a Delaunay approach. Chew's work was extended by Ruppert [14] to handle varying degrees of refinement (and thus establishing the Bern et al. optimality properties), and later by Chew also [8] to curved surfaces.

In three dimensions, no work previous to ours guaranteed bounded aspect ratio triangulations, although Chazelle and Palios [6] developed an algorithm with the best possible bound (up to a constant factor) on the cardinality of the triangulation in terms of reflex angles.

Our triangulation uses *Steiner points*, meaning that it introduces new vertices into the domain not present in the original input. Indeed, as shown by Schoenhardt, Steiner points are necessary for triangulating nonconvex polyhedra in dimensions 3 and higher. For additional background on the optimal triangulation literature, we refer the reader to the excellent surveys of Bern and Eppstein [3] and Bern and Plassmann [5]. Note that, because of the importance of mesh generation, there is a vast body literature on mesh generation algorithms. We do not attempt to survey this literature here because the majority of these papers are not concerned with mathematical quality guarantees.

The remainder of this paper is organized as follows. In Section 2 we describe the class of allowable input domains for QMG. In Section 3 and Section 4 we present a high-level description of the QMG algorithm. In Section 5–Section 7 we provide more details about the algorithm. In Section 8 and Section 9 we define aspect ratio formally and sharp angles, and establish some results about them. In Section 10–Section 17 we provide the analysis of QMG, including the proofs of the two main optimality properties mentioned above. In Section 18, we consider the asymptotic running time of QMG, and in Section 19 we briefly describe the implementation.

We remark that this paper has a companion paper [12] that describes how to triangulate a grid of uniform boxes cut by a k-affine space. The method in that paper is used as a subroutine in this paper, and we need some of the results of the analysis in that other paper for the analysis in Section 10.

Besides the QMG algorithm and its analysis, the other main contribution of this paper is a series of new bounds that apply to any possible triangulation of a polyhedral domain (see Section 9) and other results that apply to any possible bounded-aspect ratio triangulation of a polyhedral domain (see Section 16). The results in these sections act as lower bounds for proving QMG's optimality, but they would be useful for the analysis of other triangulation algorithms.

This paper, along with the companion [12] supersedes our earlier work [13]. We briefly summarize the difference between this paper and the earlier work for the reader familiar with that work. First, this work applies to d-dimensional regions for any d whereas the earlier work was limited to three dimensions. A consequence of this generalization is that we have discarded the case-based proofs used in [13] in favor of more uniform treatment here. The notion of enforcing a "balance" condition in the quadtree has been dropped. The idea of "warping" has been replaced by the approach in the companion paper, together with the "alignment" procedure described in this paper.

2 Nonconvex polyhedra

Recall that the input to our algorithm is a nonconvex polyhedron P in R^d. Mathematically, a nonconvex polyhedron is the set resulting from a finite number of union and intersection operations applied to halfspaces. We assume P is compact. We assume that P is presented via a **boundary representation**; in fact, from now on, we refer to polyhedra as "b-reps." The boundary representation of P consists of a lattice of faces: zero-dimensional faces are called vertices, one-dimensional faces edges, and the d-dimensional face is P itself. Each face of dimension 1 or higher has boundaries that are faces of one lower dimension. Thus, a brep is stored as a layered directed acyclic graph with one node for each face, and arcs to indicate the "is-a-boundary-of" relation. Nodes at level 0 (vertices) have coordinates stored with them.

Finally, we assume that P is a d-manifold with boundary to simplify our presentation, although the implementation of QMG allows many nonmanifold features such as internal boundaries.

3 Boxes

The main data structure of QMG is a box. A **box** is a d-dimensional cube embedded in an axis-parallel manner in $\mathbb{R}^d$. Our algorithm is a *quadtree*-based algorithm, meaning that it starts with a single d-cube, and then subdivides into 2^d equal-sized smaller cubes. The subdivision continues recursively.

Boxes of dimension less than d occur as separate data items. These lower-dimensional boxes are discussed in more detail in Section 7. We ignore the existence of these lower-dimensional boxes until Section 7 to allow a simplified presentation of QMG's quadtree generation in the next three sections.

Initially, there is one large d-dimensional box, called the **top box**, which contains all of P and also a neighborhood around P. This box is considered *active*. Other boxes are generated from the top box by applying one of three operations recursively. First, an active box may be *split*, meaning that it is replaced by 2^d smaller boxes each of equal size, as mentioned above. Second, a box may be *duplicated*, meaning that it is replaced by two or more boxes with the same size and position as the original box. A final operation on an active box is *protecting* it, in which case it is no longer active and no longer available for splitting or duplicating. The collection of boxes is called a *quadtree*.

The data items stored with a box are as follows. QMG stores its position and size. Because of the dyadic nature of the quadtree, the position and size are both represented exactly (as integers). As mentioned in the last paragraph, boxes are either *active* or *protected*. An active box B has stored with it its **content** which is denoted $\mathrm{co}(B)$. The definition of content is as follows. Let $\mathrm{ex}(B)$ denote a cube in $\mathbb{R}^d$ which is concentric with B but has a diameter larger by a constant factor $1+\gamma$, where γ is defined below. Note that $P \cap \mathrm{ex}(B)$ is a polyhedral region. If $P \cap \mathrm{ex}(B)$ is connected (in the topological sense), then we define $\mathrm{co}(B) = P \cap \mathrm{ex}(B)$. If $P \cap \mathrm{ex}(B)$ is not connected, then QMG makes duplicates of B, one for each component of $P \cap \mathrm{ex}(B)$, and assigns one component to each duplicate. Thus, $\mathrm{co}(B)$ is always a connected polyhedral region. More details are given in Section 5.

A protected box is always associated with a particular face F of P, and F must meet $\mathrm{ex}(B)$. Thus, a protected box has stored with it a reference to F and also a close point. The *close point* is a point in $\mathbb{R}^d$ lying in $F \cap \mathrm{ex}(B)$. The coordinates of the close points are stored in an auxiliary table, and the protected box stores an index into this table. (This is because several protected boxes can share the same close point.) The collection of close points make up the vertices of the final triangulation.

4 High level description of the quadtree generation

The mesh generation algorithm has two parts: quadtree generation and triangulation. See Fig. 1–2 for the high-level outline of quadtree generation. Triangulation is described in Section 7. Not all the terms in these figures have been defined yet.

Quadtree generation is divided into $d+1$ *phases* numbered $0, \ldots, d$. We use k throughout the paper to denote the current phase. Phase k works primarily with the k-dimensional faces of P. (Thus, in phase d we look at P itself.) Each phase is subdivided into two *stages*, the separation stage and the alignment stage. During the separation stage, active boxes are split. There is also splitting of active boxes during the alignment stage. The alignment stage also turns some active boxes into protected boxes.

5 Separation stage

In this section we describe the separation stage of phase k in more detail. At the start of the phase there is a list of active boxes I_k and the (initially

/* Quadtree generation */.
Initialize $I_0 := \{\text{top\_box}\}$.
Initialize $J := \{\}$.
for $k := 0, \ldots, d$ do
    Initialize $I_{k+1} := \{\}$.
    Initialize $O_F := \{\}$ for each k-dimensional P-face F.
    /* Phase k separation stage. */
    while I_k is nonempty do
        Remove an active box B from I_k.
        if B is crowded or too big for the size function then
            Split B into $B_1, \ldots, B_{2^d}$; duplicate as necessary.
            Delete B_i's with empty content.
            Put remaining B_i's into I_k.
        elseif $\mathrm{co}(B)$ contains a (necessarily unique) k-face F of P then
            $O_F := O_F \cup \{B\}$.
        else
            $I_{k+1} := I_{k+1} \cup \{B\}$.
        end if
    end while

Figure 1: High-level description of QMG's quadtree generation (continued in Fig. 2).

/* **Phase k alignment stage.** */
for each k-dimensional P-face F do
  while O_F is nonempty do
    Remove the highest-precedence box B from O_F.
    Find the highest priority subface B' of B that is close to F.
    if B has no such close subface then
      $I_{k+1} := I_{k+1} \cup \{B\}$.
    elseif the alignment condition is satisfied for B then
      Protect B; its associated P-face is F.
      Find the close point on F for B (near B').
      $J := J \cup \{B\}$.
    else
      Split B into $B_1, \ldots, B_{2^d}$; duplicate as necessary.
      Delete the B_i's with no content.
      Put remaining B_i's into O_F.
    end if
  end while
end for
end for /* **end of k loop** */

Figure 2: High-level description of QMG's quadtree generation (continued from Fig. 1).

empty) idle list I_{k+1}. During the phase we repeatedly remove one active box, say B, from I_k and test it for crowdedness (defined below). If B is crowded, then it is split. Let us use the term "children" to denote the boxes on the next deeper level resulting from the split. All the children with a nonempty content are inserted back into I_k. Box B itself is deleted, and the children with empty content are deleted. (Boxes with empty content that arise during splitting for alignment are also deleted.) On the other hand, if B is not crowded, then we check whether it has a P-face of dimension k, say F, in its content. If so, the box is transferred to orbit O_F. If not, the box is transferred to the idle list. In this manner I_k is eventually emptied.

We now explain the terms "content" and "crowdedness." First, we define $\mathrm{ex}(B)$ for an active box B to be a d-dimensional cube in $\mathbb{R}^d$ concentric with B but expanded in each dimension by a multiplicative factor $1 + \gamma$. Parameter γ must satisfy $\gamma \geq \epsilon_{0,F}$ for each P-face F, where $\epsilon_{0,F}$ is the tolerance for alignment described in Section 6. For instance, $\gamma = 0.5$ is acceptable.

The **content** of an active box B is a b-rep and is typically $P \cap \mathrm{ex}(B)$. However, if $P \cap \mathrm{ex}(B)$ has more than one connected component, we identify the components of $P \cap \mathrm{ex}(B)$, say $C_1, \ldots, C_p$ and we replace B with p copies of itself, say $B_1, \ldots, B_p$. Then we define $\mathrm{co}(B_i) = C_i$ for $i = 1, \ldots, p$.

Say B is split, and say B' is one of the child boxes. We compute $\mathrm{co}(B')$ by intersecting $\mathrm{co}(B)$ with $\mathrm{ex}(B')$. (Notice that our definition of $\mathrm{ex}(B)$ guarantees that $\mathrm{ex}(B')$ is a proper subset of $\mathrm{ex}(B)$.) In particular, we do not compute $\mathrm{co}(B')$ by intersecting the original b-rep P with $\mathrm{ex}(B')$. This is because this latter approach could reintroduce connected components that were duplicated into a box different from B at some earlier level of splitting.

We say that a box B is **crowded** if (1) $\mathrm{co}(B)$ meets any P-face of dimension $k-1$ or less, or (2) $\mathrm{co}(B)$ meets a P-face F of dimension k, and $\mathrm{co}(B)$ meets another P-face G that is not a superface of F.

Thus, a box B is not crowded during phase k if either (1) $\mathrm{co}(B)$ does not meet any P-faces of dimension k or lower or (2) $\mathrm{co}(B)$ meets exactly one P-face F of dimension k, no P-faces of dimension less than k, and every P-face of dimension higher than k in $\mathrm{co}(B)$ is a superface of F. A box that is not crowded is either transferred to I_{k+1} if (1) holds, or it is transferred to O_F if (2) holds. Some examples of crowdedness are given in Fig. 3.

Another rule used in the separation stage is that we split boxes if their side-length is greater than the user-specified mesh refinement function that was mentioned in the introduction. We do not say any more about this here, since the analysis in subsequent sections does not involve a user-specified

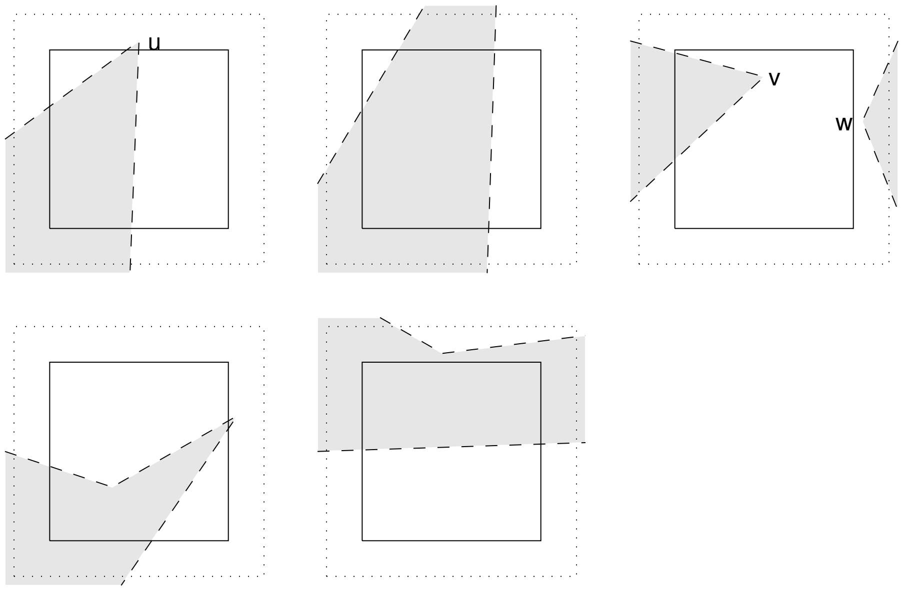

Figure 3: Suppose we are in the separation stage of the phase 0 in the case $d = 2$. In the above figures, solid lines indicate boxes, dotted lines indicate $\mathrm{ex}(B)$ for these boxes, dashed lines indicate the boundary of P, and shading represents the interior of P. All boxes in the top row are uncrowded. The first box would be placed into $O_{\boldsymbol{u}}$. The second box would be placed into I_1. The third box in the top row must be duplicated, and then one duplicate would go into $O_{\boldsymbol{v}}$ and the other into $O_{\boldsymbol{w}}$. Both boxes in the bottom row are crowded and must be split.

mesh refinement function.

The reader may notice that there appears to be a potential infinite loop: if an active box B in phase k has a P-face of dimension $k-1$ or less in its interior, then this will cause an infinite recursion of splitting because there will always be a crowded subbox. Fortunately, this situation can never occur. The reason is that a box whose interior meets a P-face of dimension $k-1$ or less would have had a close subface identified in an earlier phase and would have become protected, or would have been crowded in an earlier phase. (See the next section for a description of close subfaces.) Therefore, it could never end up in I_k. It is possible however, for a box in I_k to have a face of dimension $k-1$ or less inside $\mathrm{ex}(B)$ but outside B. This can happen because, in the previous phase, a $(k-1)$-face F could lie in the content of B and yet not be close enough to come close to a subface of B. In this case the box will be split until the d-cubes $\mathrm{ex}(B)$ have shrunk enough that they do not meet the low-dimensional P-face. The number of times that a box can be split is analyzed in subsequent sections.

The computation of $\mathrm{co}(B)$ (that is, computing the geometric intersection $\mathrm{co}(B') \cap \mathrm{ex}(B)$, where B' is the parent of B, and then checking whether this intersection is connected) is among the most computationally intensive tasks of QMG. We carry out the search for connected components with a ray-shooting algorithm that we do not describe here. The worst-case running time of this ray-shooting algorithm is $O(n^2)$, where n is the total geometric complexity of $\mathrm{co}(B)$ (i.e., the total number of boundary faces), but in practice the running time will usually be closer to $O(n)$.

In the case $d=2$, it is possible to find connected components of $\mathrm{co}(B)$ via a plane sweep in $O(n \log n)$ operations. In the case $d=3$, an $O(n \log n)$ plane sweep can also be used provided that P is preprocessed with $O(N^2)$ preprocessing steps, where N is the combinatorial complexity of the original P. This efficient algorithm for $d=3$ is described in our earlier paper [13]. We have not implemented a plane-sweep procedure for either $d=2$ or $d=3$.

6 Alignment

In this section we describe the alignment stage. Recall that the alignment stage processes each orbit independently. For this section, assume we are in phase k and are processing orbit O_F of P-face F whose dimension is k.

First, a sequence of parameters

$$0 < \epsilon_{d-k-1,F} < \epsilon_{d-k-2,F} < \cdots < \epsilon_{0,F} < 0.5$$

is chosen for F. The method for choosing these parameters is described in in [12], but must be slightly modified to take into account the containment relationship between P-faces of different dimensions. These parameters have upper and lower bounds depending only on d and k.

We now process boxes in O_F in a precedence order to be described below. Let B be the high-precedence box in the orbit. Let B' be any subface of B. We construct the ∞-norm neighborhood of radius $\epsilon_{r,F}$ around B', denoted $N(B')$, where r stands for the dimension of B'. Thus, this neighborhood is an axis-parallel parallelepiped (which could be degenerate if $\epsilon_{r,F} = 0$). If F passes through $N(B')$, then F is said to be **close** to B'. The **close subface** of B is the box subface of lowest dimension that is close to F. If there is a tie (i.e., there are several faces of the same lowest dimension all close to F), then we break the tie with a priority rule, which is described below. A box with no close subface is transferred to I_{k+1}.

Because $\gamma \geq \epsilon_{0,F}$, if F is close to B (i.e., if B has a close subface), then F must pass though $\text{ex}(B)$. Thus, we can check whether B has a subfce close to F by examining $\text{co}(B)$. Indeed, it is important that we query $\text{co}(B)$ rather than the original P, because it might be difficult to determine from queries on P whether the P-face in question is associated with B or with a duplicate of B.

Next we claim a partial converse: if F meets $\overline{\text{ex}}(B)$, then B has a subface close to F. We define $\overline{\text{ex}}(B)$ to be a cube concentric with B and expanded by $\epsilon_{d-k-1,F}$ in each dimension; thus $B \subset \overline{\text{ex}}(B) \subset \text{ex}(B)$. It follows from Lemmas 1 and 2 of [12] that if any P-face F passes through $\overline{\text{ex}}(B)$, then B has a subface close to F. The cube $\overline{\text{ex}}(B)$ is not used in our algorithm, but it plays a role in the analysis below.

Once every box in the orbit has chosen its close subface, we now test the **alignment condition**. The alignment condition is as follows. Define the **extended orbit** of F to be O_F united with protected boxes from earlier phases that are associated with proper subfaces of F. For every active box B in the orbit, the close subface of B must be completely covered by boxes in the extended orbit (either active or protected) that are the same size or larger. For an example of the alignment condition, see Fig. 4. We provide motivation for the alignment condition in Section 7

Let us now comment further on the alignment condition. First, we have to explain what is meant by "completely covered." We say that a box subface B' is completely covered by some collection of boxes $\{B_1, \ldots, B_n\}$ provided that for any point $\boldsymbol{p}$ in the relative interior of B', there exists an open neighborhood N of $\boldsymbol{p}$ such that $N \subset B_1 \cup \cdots \cup B_n$. (Note that if B' is a 0-dimensional box subface, i.e., a vertex, then its relative interior is B' itself.)

With these definitions of "extended orbit" and "completely covered," we can now state the priority rule for choosing a close subface. Recall that the close subface of B is the box subface of lowest dimension close to B. Let l be the dimension of this subface. If there is a tie (i.e., there is more than one face of B of the dimension l close to F), then we favor the subfaces that are completely covered by boxes the same size or larger in the extended orbit, i.e., those close subfaces of dimension l for which the alignment condition holds. If there is still a tie, then we use a lexicographic tie-breaking rule.

Recall that boxes can get duplicated during the separation stage, and thus several active boxes can cover the same geometric region in $\mathbb{R}^d$. We claim that two boxes with overlapping geometric regions in $\mathbb{R}^d$ cannot end up in the same orbit. (This fact simplifies the sorting necessary to check the alignment condition.) The reason is as follows. Suppose B and B' are two boxes whose interiors have a common point in $\mathbb{R}^d$, and suppose both $\mathrm{co}(B)$ and $\mathrm{co}(B')$ contain a point of P-face F. By the tree-nature of the quadtree, two boxes that share a common interior point must have the property that one is contained in the other.

We claim that $\mathrm{co}(B)$ and $\mathrm{co}(B')$ must both meet a proper subface of F. If not, then $\mathrm{co}(B)$ would have to contain the intersection of $\mathrm{ex}(B)$ with the entire affine hull of F (because no boundaries of F are in $\mathrm{co}(b)$). Similarly, $\mathrm{co}(B')$ would also contain the intersection $\mathrm{ex}(B')$ with the hull of F. (See Section 8 for definition of "affine hull" and other mathematical terminology.) Since one box contains the other, this means that one box contains points from F that the other box also contains. But then there could not be two distinct connected components of P in $\mathrm{co}(B)$ and $\mathrm{co}(B')$, so duplication would not have taken place.

Thus, $\mathrm{co}(B), \mathrm{co}(B')$ each contain proper subfaces of F. But in this case, the boxes could not end up in O_F (i.e., if they were still active in phase $\dim(F)$, they would be crowded).

As mentioned earlier, a box is protected if the alignment condition holds for its close subface. We make the following claim: if the alignment condition holds for B at the time it is protected, then the condition continues to hold for

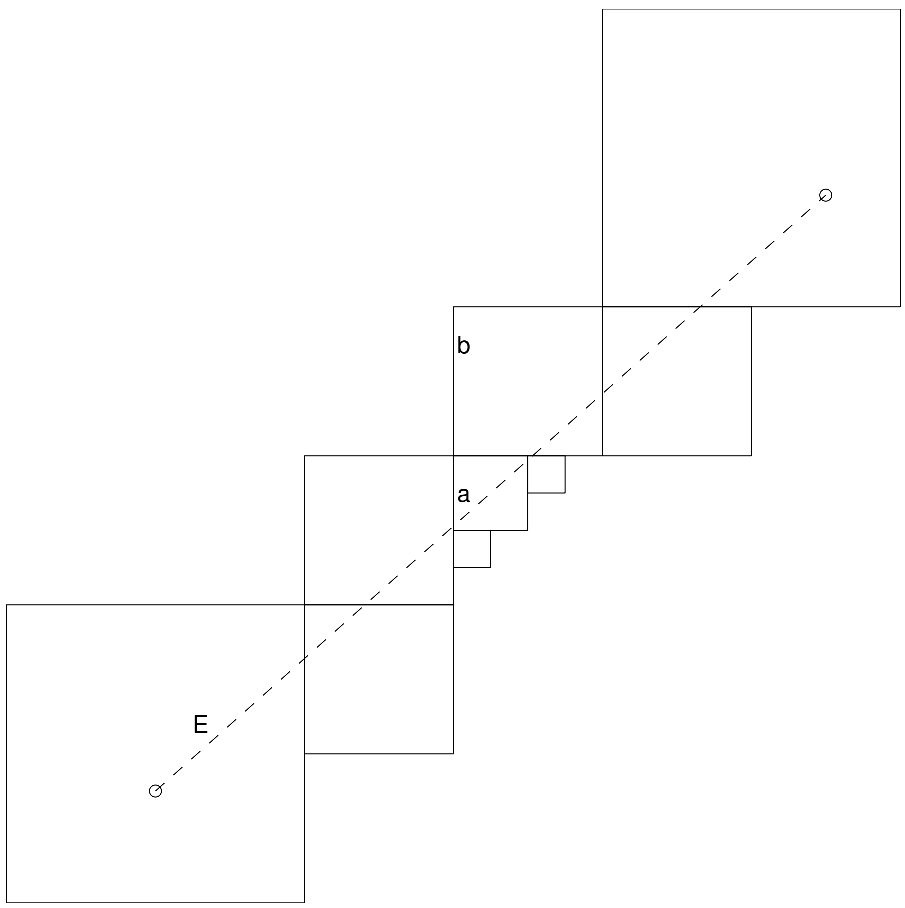

Figure 4: In this figure we illustrate the alignment condition in the case $d = 2, k = 1$. The boxes in this figure are the extended orbit of a P-edge E, which is the dashed line. The two large boxes at the ends are protected boxes for the endpoints of E, protected from phase 0. In this figure, box a must be split because the alignment condition does not hold for this box. Its close subface, which could be either its lower left-hand corner or upper right hand corner, is contained by another box smaller than a. All other boxes satisfy the alignment condition. For example, box b does not have to be split; its close subface could be either its bottom edge or right edge. The right edge will have higher priority since the alignment condition holds for that edge.

the remainder of the algorithm. In other words, the following situation cannot occur: A box B with close subface B' is deemed to satisfy the alignment condition and protected. Later a neighboring box $\bar{B}$ also containing B' as a subface gets split because the alignment condition does not hold for $\bar{B}$, thus causing the alignment condition to be violated for B.

In order to prove the claim in the last paragraph, we must describe the order in which QMG processes the boxes in an orbit O_F. "Process" means that QMG determines whether the box satisfies the alignment condition; if so, then protect it, and if not, then split it. The correct order is to start with the largest boxes in the orbit, working down to the smallest. Within the set of boxes of the same size, we process those with the lowest-dimensional close subfaces first, working towards highest-dimensional close subfaces.

We claim that this order assures the claim that if the alignment condition holds for a box B at the time it is processed, then the alignment condition holds for B for the remainder of the algorithm. Suppose we are at the step when B is processed and the alignment condition is satisfied. Let B' be the close subface of B, and let l be the dimension of B'. Let $B_1, \ldots, B_n$ be the boxes in the extended orbit that cover B'. Some of $B_1, \ldots, B_n$ will be larger than B' and hence already protected. Protected boxes are not split again, so they will continue to cover B' for the rest of the algorithm. Consider a box B_i that is the same size as B. If B_i has a close face of dimension less than l, then B_i is already protected (because we process boxes with lower-dimensional close faces first). The dimension of the close face of B_i cannot be greater than l, because B' is a subface of B_i and has higher priority than any subface of B_i of dimension $l+1$ or more. Therefore, the only remaining possibility is that B_i is the same size as B and that the close subface of B_i has dimension exactly l. But then this subface, if it is not B', must also be completely covered by boxes in the orbit because otherwise B' would have higher priority. (Recall that faces that are completely covered have higher priority.) So we see that B_i will become protected as well and cannot be split.

When a box B is protected, as mentioned above, we have identified a close subface B' of B. This subface has the property that F passes through an ∞-norm neighborhood of B'. We now select a point lying on F in this neighborhood (see our other paper [12] for more details on selecting the close point). The rule used for choosing the close point has the property that any other box B'' that is the same size as B and also has B' as its close subface will choose the same close point. Thus, several adjacent boxes that are in the same orbit and are the same size might share a close point.

The alignment stage continues until there are no boxes left to process; every box is either protected or has been moved to the idle list. When O_F is empty, the alignment moves onto a different orbit. Once the orbits of all dimension-k faces of P are empty, the phase is over.

7 Triangulation

After phase d of quadtree generation, QMG triangulates the quadtree. In the triangulation procedure, the collection of protected boxes is triangulated into a simplicial complex.

In order to describe the triangulation procedure, we must first bring lower-dimensional boxes into the picture. In this section, we revisit some of the concepts from earlier sections and revise some of the algorithm steps to take into account lower dimensional boxes. The lower dimensional boxes serve two purposes: first, they simplify the data structures needed for checking the alignment condition and second, they serve as the basis for generating the final triangulation.

In QMG, boxes can have dimension 0 up to d. Initially, there is only one active box of dimension d, namely, the top box. Lower dimensional boxes get created each time a box is protected by QMG. At the moment an i-dimensional active box B is changed from active to protected during the alignment stage for orbit O_F, all its faces of dimension $i-1$ are launched as new active boxes (there are $2i$ such faces). Let B' be one of these new active boxes. It is dealt with in a manner analogous to the way QMG handles subboxes after a split. We compute $\mathrm{co}(B')$ as the intersection $\mathrm{co}(B)\cap\mathrm{ex}(B')$. We determine the disposition of B' using the same rules as before: If $\mathrm{co}(B')$ is empty, then we delete B'. If $\mathrm{co}(B)\cap\mathrm{ex}(B')$ has more than connected component, then we duplicate B'. If $\mathrm{co}(B')$ does not meet F (and hence meets only P faces of dimension $k+1$ and higher) then we place B' in I_{k+1}. If $\mathrm{co}(B')$ meets F, then we place B' in O_F.

It is possible that B' will also become immediately protected; this happens, for instance, when the close subface of B is also a subface of B'.

The same operations are performed on an i-dimensional box as on a full-dimensional box: such a box can be tested for crowdedness, split for separation, protected and so on. When an i-dimensional box is split, 2^i new subboxes are created. If B is i-dimensional, then it is said to **extend over**

i of the possible d coordinate axes, and it is **flat** over the remaining $d - i$ coordinate axes.

The definition of $\mathrm{ex}(B)$ for a box of dimension less than d is as follows. Every box B has associated with it a number called its *size*, which we denote $\mathrm{size}(B)$ and which is the side-length of B in a dimension over which it extends. The size of every box is equal to the size of the top box multiplied by a factor 2^{-p}, where p is the number of times the top box was split to reach this box. If B is a box, then $\mathrm{ex}(B)$ is an axis-parallel full-dimensional rectangle in $\mathbb{R}^d$ centered at the center of B, with side-lengths $(1+\gamma)\,\mathrm{size}(B)$ for axes over which B extends, and side-length $\gamma\,\mathrm{size}(B)$ for the axes in which B is flat. This choice ensures that all properties of $\mathrm{ex}(B)$ asserted earlier are still valid, namely, if B has a subface close to F, then the subface passes through $\mathrm{ex}(B)$. Also, if B is split, then $\mathrm{ex}(B_j)$ for each subbox B_j is contained in $\mathrm{ex}(B)$. Finally, if subfaces of B are launched as new active boxes when B is protected, then each new subface B_j also satisfies $\mathrm{ex}(B_j) \subset \mathrm{ex}(B)$. Note that for consistency, even zero-dimensional boxes must have a size. When a zero-dimensional box B is split, there is only one child, but splitting still has significance because the size is halved, which diminishes $\mathrm{ex}(B)$ and therefore $\mathrm{co}(B)$.

Earlier, when describing the alignment condition, we introduced the terms "extended orbit" and "completely cover." Recall that we defined extended orbit to be the union of the orbit O_F of a face F united with the protected boxes for all proper subfaces of F from previous phases. In fact, QMG never forms extended orbits; instead, the lower-dimensional active box faces of protected boxes act as proxies for the protected boxes. A system of weights is used to determine the complete coverage condition. In particular, every active box in QMG stores a weight associated with each of its subfaces. Thus, an i dimensional active box has 3^i weights stored with it. Each weight is a number between 0 and 1 that indicates what fraction of the subface is "owned" by that active box. Initially, the top box owns all of its subfaces. When a box is split, the weights are divided up among children. We omit the details of how the weights get split up, but the upshot is that QMG can test whether a box subface is completely covered by boxes in its orbit by adding up the weights associated with that subface contributed by all the boxes containing it; complete coverage is indicated by a weight sum of 1.0.

Although many details are omitted, we do mention one key point that reduces the amount of searching and sorting in QMG. When testing the complete coverage rule, it is necessary to look only at boxes of a single size.

This means that the complete coverage condition can be tested with a simple hash-table. Consider the example in Fig. 5.

It can be shown that these more complicated rules introduced in this section are equivalent to the definitions in the previous sections in the following sense. For a given input P, the quadtree generation procedure produces the same sequence of full-dimensional boxes whether we follow the rules of this section or preceding sections.

The protected boxes are linked together by pointers; in particular, a protected box of dimension i has pointers to all the protected boxes of dimension $i+1$ of which it is a subface. This data structure serves as the basis for triangulation.

The triangulation algorithm is based on our other paper [12] and is as follows. Let a **chain** be a sequence of nested boxes $B_0, \ldots, B_d$ such that the dimension of B_i is i. "Nested" means that, for each i, B_i is a face or subset of a face of B_{i+1}. Let $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_d$ be the close points of $B_0, \ldots, B_d$. Then the simplex whose vertices are $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_d$ is put in the triangulation. Thus, the triangulation has one simplex for each chain. The only exception is when a close point is repeated in this chain; in this case, the simplex is said to be **null** and is not included in the triangulation. QMG enumerates all possible chains with a stack-based search algorithm. An example of the triangulation algorithm is presented in Fig. 6.

We can now explain the importance of the alignment condition in Fig. 7. As is seen from the figure, if the alignment condition were not enforced, then the triangulation algorithm described in the previous paragraph would be invalid.

8 Aspect ratio

In our analysis of QMG, which begins in Section 10, we demonstrate two optimality properties: the triangulation generated by QMG has optimal aspect ratio, up to a certain factor, and also optimal cardinality (compared to all other bounded-aspect ratio triangulations), up to a certain factor. Before demonstrating these properties, we must provide the definitions aspect ratio, sharp angle, and so. This mathematical background material is the topic of this section and the next section.

First, we provide some standard definitions from linear algebra. An *affine set* X is the solution to a system of linear equations, that is, $X = \{\boldsymbol{x} \in \mathbb{R}^d :$

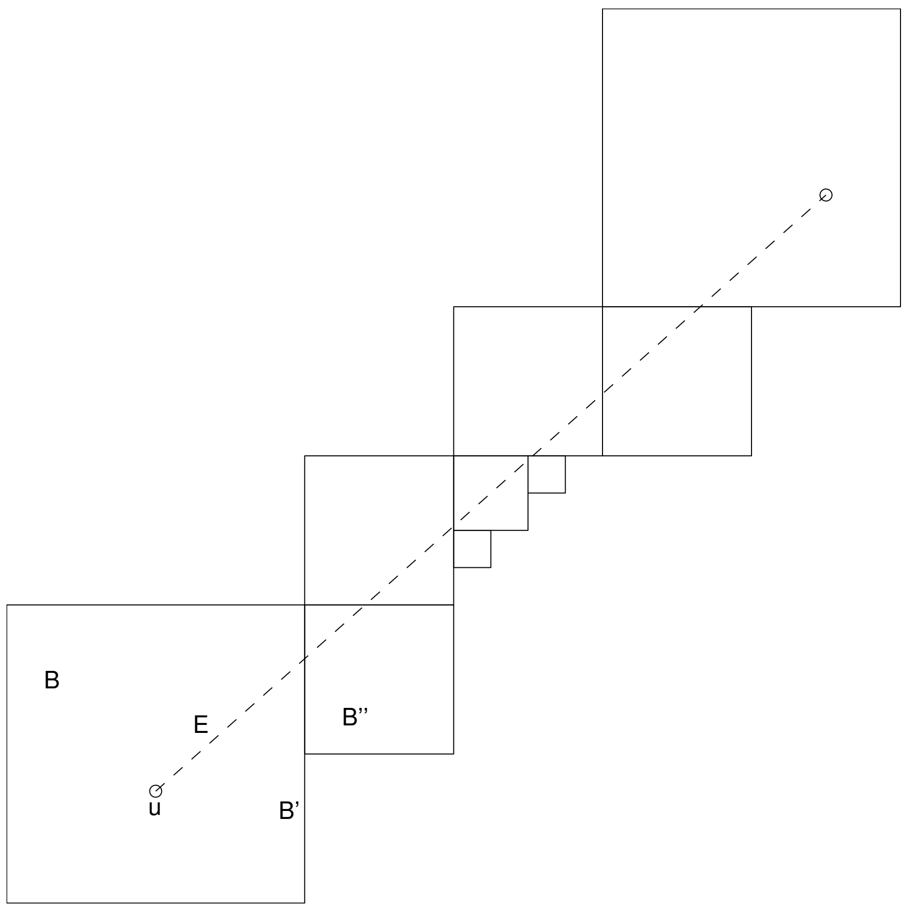

Figure 5: This figure revisits the alignment condition taking into account lower dimensional boxes. Focus on the lower left-hand box B, which is protected in phase 0 for the vertex u of P in its interior. The right edge of this box—call it B'—becomes a new active box; it is uncrowded and its associated P-face is E (E is the dashed segment), and its close subface is the whole box B'. During the alignment phase 1, this 1-dimensional box is split because the weight of B' associated with B' itself is only 0.5. Once B' is split in half, its lower half no longer meets E and hence is placed into I_2. Its upper half together with the full dimensional active box labeled B'' completely cover the upper half of B' so the alignment condition is satisfied.

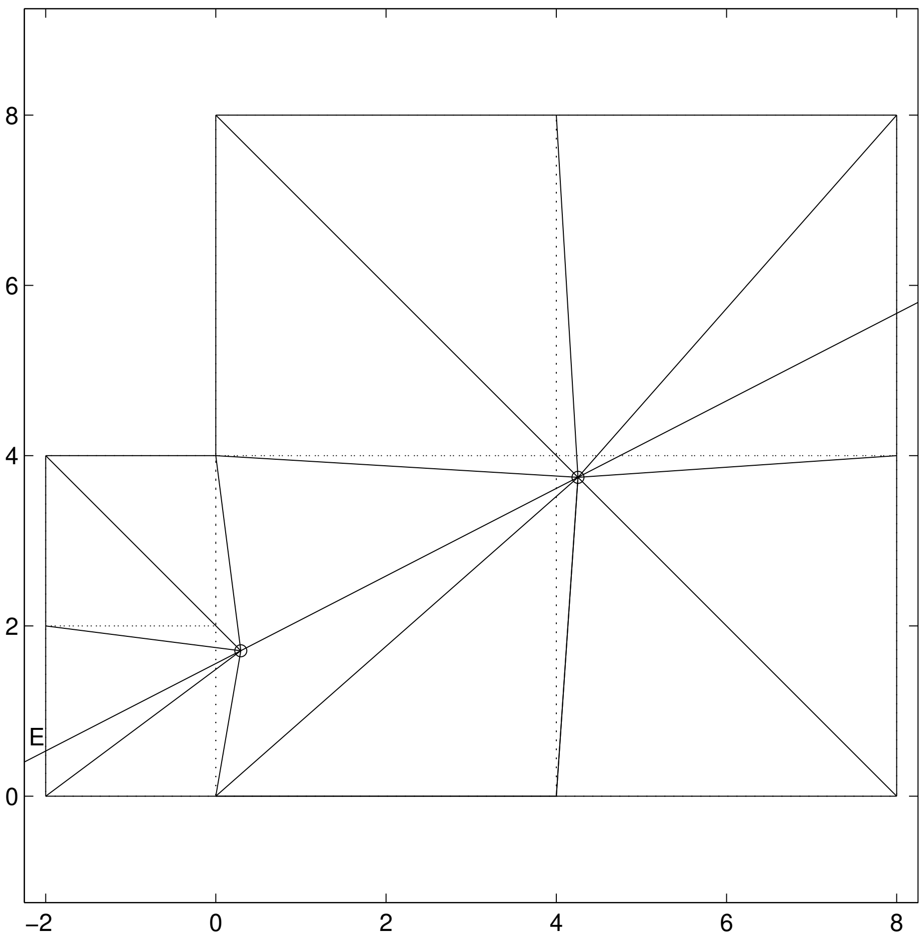

Figure 6: An example of the triangulation procedure in the case $d = 2$. All six full-dimensional protected boxes in this figure are associated with P-edge E. The dotted lines are boundaries of boxes that are not present in the final triangulation. The solid segments are all part of the triangulation. The close point for the two small boxes on the left is the point near $(0, 2)$ marked in the figure. An example of a non-null chain in this figure would be starting from the vertex $(0, 0)$ (whose close point is at $(0, 0)$ and is associated with the two-dimensional face P itself), then the edge $\{0\} \times [0, 2]$ containing it, whose close point is the marked point on E near $(0, 2)$, and finally the box $[0, 4] \times [0, 4]$, whose close point is the marked point near $(4, 4)$ on E. An example of a null chain would be the vertex at $(4, 0)$ (whose closepoint is at $(4, 0)$), the edge $\{4\} \times [0, 4]$ (whose close point is on E near $(4, 4)$), and finally the box $[4, 8] \times [0, 4]$, which has the same close point near $(4, 4)$. This figure shows a triangulation on both sides of E for illustrative purposes, although if E were a boundary edge, in fact only one side of E would be triangulated. As mentioned in the introduction, however, the actual implementation of QMG allows internal boundaries so the above situation of triangulating both sides of an edge does occur with QMG.

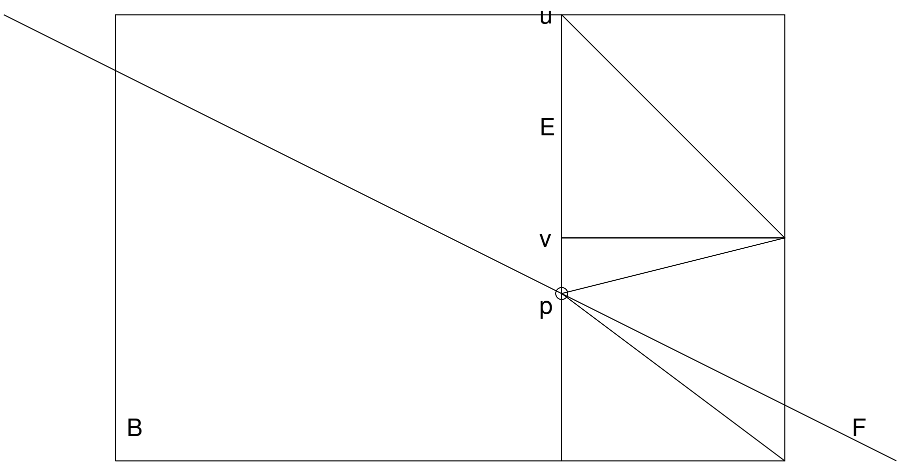

Figure 7: The above figure shows an inconsistency that would result during triangulation if the alignment condition were not enforced. In particular, the close point of box B is point $\boldsymbol{p}$, which lies on P-face F and also on the right edge of box B. Thus, the close face for B is this edge. The alignment condition does not hold, i.e., the close face of B is covered by smaller boxes on the right. Note that we obtain an illegal triangulation. For instance, the close point of edge E is one of its endpoints, say endpoint $\boldsymbol{v}$. Then the chain consisting of $\boldsymbol{u}$, then E (whose close point is $\boldsymbol{v}$), and then B (whose close point is $\boldsymbol{p}$) is a flat simplex, i.e., a triangle whose three points are collinear. Such a simplex has infinite aspect ratio and must not be allowed in a triangulation. Note that a degenerate simplex like this is not a null simplex defined in Section 7: this degenerate simplex does not have any repeated vertex and hence it cannot be legally dropped from the triangulation.

$A\boldsymbol{x} = \boldsymbol{b}\}$ for some $m \times d$ matrix A with linearly independent rows and some m-vector $\boldsymbol{b}$. The *dimension* of this affine set is $d - m$. Let Y be any subset of $\mathbb{R}^d$. The *affine hull* of Y is defined to be lowest-dimensional affine set that contains Y and is denoted $\mathrm{aff}(Y)$. It can be shown that $\mathrm{aff}(Y)$ is uniquely determined by this definition. In particular, it can be shown that $\mathrm{aff}(Y)$ is the set of all points that can be written in the form $\alpha_1 \boldsymbol{y}_1 + \cdots + \alpha_s \boldsymbol{y}_s$, where s is an arbitrary positive integer, $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_s \in Y$, and $\alpha_1, \ldots, \alpha_s$ is an arbitrary sequence of real number that add up to 1. Let F be a face of P. Since P is polyhedral, $\mathrm{aff}(F)$ and F have the same dimension.

We now define define aspect ratio.

Definition. *Let T be a d-simplex in $\mathbb{R}^d$ with vertices $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_d$. Then the* **altitude** *of T at $\boldsymbol{v}_i$ is defined to be* $\mathrm{dist}(\boldsymbol{v}_i, \mathrm{aff}(\boldsymbol{v}_0, \cdots, \boldsymbol{v}_{i-1}, \boldsymbol{v}_{i+1}, \cdots, \boldsymbol{v}_d))$. *The* **minimum altitude** *of T, denoted* $\mathrm{minalt}(T)$*, is the minimum altitude over all choices of $\boldsymbol{v}_i$ for $i = 0, \ldots, d$.*

Definition. *The* **aspect ratio** *of a simplex T is defined to be*

$$\mathrm{asp}(T) = \mathrm{maxside}(T) / \mathrm{minalt}(T).$$

Thus, the aspect ratio is always at least 1, and large aspect ratios indicate poor quality elements.

In the remainder of this section, we characterize aspect ratio in terms of matrix norms. Given a d-simplex T, we define its associated matrix M_T to be the $d \times d$ matrix whose ith column, $i = 1, \ldots, d$, is $\boldsymbol{v}_i - \boldsymbol{v}_0$. Thus, M_T depends on the numbering of the vertices, and in particular, $\boldsymbol{v}_0$ plays a distinguished role. However, we note the following: if we define M_T' according to a different numbering of the vertices, then the columns of M_T' can be obtained from the columns of M_T by subtracting pairs of columns in M_T and then permuting. In linear algebra terms, there exists a $d \times d$ matrix L all of whose entries are zeros except for possibly one '1' and one '-1' in each column such that $M_T' = M_T L$, and such that L^{-1} has the same properties (all zeros except for possibly one '1' and one '-1' per column).

The following two results hold for any numbering. These lemmas use the following well-known linear algebra fact. The norm of a $d \times d$ matrix is bounded above and below by constant multiples (where the constant depends on d) of the maximum norm among its columns, and also above and below by constant multiples of the maximum norm among its rows. In the remainder of this paper c_d denotes a constant depending only on d that may change from formula to formula.

Lemma 1 *Let σ denote* $\operatorname{maxside}(T)$. *Then*

$$c_d\sigma \leq \|M_T\| \leq C_d\sigma$$

where c_d, C_d are two constants depending only on d.

Proof. There are two cases, depending on whether σ is the length of a side adjacent to $\boldsymbol{v}_0$ or not. In the first case, say that $\sigma = \|\boldsymbol{v}_1 - \boldsymbol{v}_0\|$. Then both inequalities are easy because σ is the norm of the first column of M_T, and all the other columns of M_T have norm bounded above by σ.

The other case is that σ is the length of a side not adjacent to $\boldsymbol{v}_0$. Then we can reduce to the first case by renumbering the vertices and noting that the norms of the transformation matrices mentioned above, $\|L\|$ and $\|L^{-1}\|$, are bounded above by constants depending only on d. ∎

Lemma 2 *Let $\mu = \operatorname{minalt}(T)$. Then*

$$c_d/\mu \leq \|M_T^{-1}\| \leq C_d/\mu$$

where c_d, C_d are two constants depending only on d (not necessarily the same constants as in the previous lemma).

Proof. Let $\boldsymbol{u}^{\mathrm{T}}$ be the ith row of M_T^{-1}. Then $M_T^{\mathrm{T}}\boldsymbol{u}$ is a column of the the identity matrix. (The superscript T denotes transpose. The subscript T indicates the association of M with simplex T.) Geometrically, this means that $\boldsymbol{u}$ is orthogonal to $d-1$ columns of M_T; in particular, $\boldsymbol{u}$ is orthogonal to the plane $\operatorname{aff}(\boldsymbol{v}_0, \ldots, \boldsymbol{v}_{i-1}, \boldsymbol{v}_{i+1}, \ldots, \boldsymbol{v}_d)$. Thus, $\boldsymbol{u}$ is parallel to the altitude from vertex i. Its length is chosen so that its inner product with $\boldsymbol{v}_i - \boldsymbol{v}_0$ is 1, i.e. its inner product with the true altitude vector is 1. Thus, the ith row of M_T^{-1} is parallel to the altitude vector from $\boldsymbol{v}_i$ but is scaled so that its length is the reciprocal of the altitude.

Then we see that the rows of M_T^{-1} have lengths equal to reciprocals of altitudes from $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_d$, with the shortest altitude being the reciprocal of the norm of the largest row. This proves the lemma, provided that the min altitude is not from $\boldsymbol{v}_0$. The case when the min altitude is adjacent to $\boldsymbol{v}_0$ is handled by renumbering as in the previous proof. ∎

We conclude from these two lemmas that $\operatorname{asp}(T)$ is within a constant factor of $\|M_T\| \cdot \|M_T^{-1}\|$, that is, the condition number $\kappa(M_T)$. Combining these lemmas with the Hadamard inequality yields the following well-known result:

$$c_d \operatorname{minalt}(T)^d \leq \operatorname{vol}(T) \leq C_d \operatorname{maxside}(T)^d. \tag{1}$$

9 Angles and PL paths

In the previous section, we defined "aspect ratio." It turns out that we can show that QMG produces triangulations whose aspect ratio is bounded above in terms of the sharpest angle of the input domain P. In this section, we provide the definition of "sharpest angle" and a theorem stating that *any possible* triangulation of P has aspect ratio bounded below in terms of the sharpest angle. Thus, the theorem in this section is the lower bound necessary to prove that the QMG triangulation is optimal.

Let $\boldsymbol{x}, \boldsymbol{y}$ be two points in P. A **piecewise linear** path Π from $\boldsymbol{x}$ to $\boldsymbol{y}$ is a path composed of a finite number of line segments. The endpoints of the segments are the **breakpoints** of Π. For the rest of the paper, we use "PL" to stand for "piecewise linear." Suppose that $\boldsymbol{x} \in F$ and $\boldsymbol{y} \in G$, where F and G are two faces of P. We will say that Π is **contractible** if there exists a point $\boldsymbol{z}$ such that the segment $\boldsymbol{xz}$ lies in F, the segment $\boldsymbol{yz}$ lies in G, and for all $\boldsymbol{v} \in \Pi$, the segment $\boldsymbol{vz}$ lies in P. Note that this definition forces $\boldsymbol{z}$ to lie in both F and G. Thus, a necessary condition for contractibility is that F and G have a nonempty common subface.

Note that we should really apply this term "contractible" to a triplet (Π, F, G), since the definition depends on the specification of F and G as well as on the path Π. When we use the term, the choice of F and G will be understood from context. The opposite of contractible is **incontractible**.

Let $\mathcal{T}$ be an arbitrary triangulation of P (not necessarily the triangulation produced by QMG). If a path Π is contractible, then we can obtain a lower bound the aspect ratios of simplices of $\mathcal{T}$ that meet Π. On the other hand, if Π is incontractible, then we can obtain an upper bound on the minimum altitude of simplices that meet Π. The remainder of this section is devoted to stating and proving these two results.

We start with the definition of the "angle" between two P-faces F and G, which is defined by contractible paths.

Definition. *Let F, G be two faces of P, and suppose $\boldsymbol{x} \in F$ and $\boldsymbol{y} \in G$. Let* Π *be a contractible PL path from $\boldsymbol{x}$ to $\boldsymbol{y}$. Let A be the affine set* $\operatorname{aff}(F) \cap \operatorname{aff}(G)$*. The* **angle** *determined by* $(\boldsymbol{x}, \boldsymbol{y}, F, G, \Pi)$ *is*

$$\theta = \frac{\operatorname{lth}(\Pi)}{\min(\operatorname{dist}(\boldsymbol{x}, A), \operatorname{dist}(\boldsymbol{y}, A))}. \tag{2}$$

We say that the **sharpest angle** *formed by F and G in the infimum of*

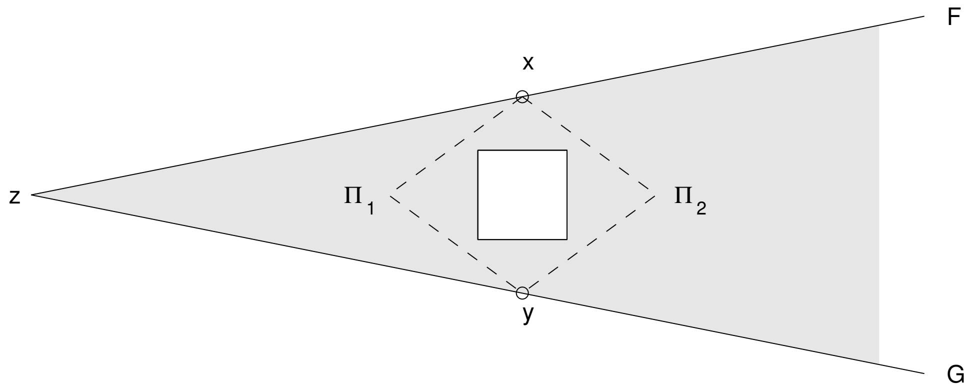

Figure 8: In the $d = 2$ case, F and G are two boundary segments meeting at boundary vertex $\boldsymbol{z}$. The polygon has a square hole in it: the interior of P is shaded. The PL path Π_1 connects a point $\boldsymbol{x} \in F$ to $\boldsymbol{y} \in G$. This path is contractible to $\boldsymbol{z}$. On the other hand, path Π_2 is incontractible.

(2) *over all contractible paths from F to G (assuming that at least one contractible path from F to G exists). Finally, we say that the sharpest angle in P is the infimum over all angles.*

In the denominator of (2), "dist" denotes ordinary Euclidean distance. Note that $\boldsymbol{z}$, the base of the contraction, does not appear in (2). An example of a contractible path is given in Fig. 8. We are concerned about the case when this angle is small, so large angles have no significance. Thus, the degenerate case when the denominator of (2) is 0 (which happens, for instance, if F is a superface of G or vice versa) is not relevant for our analysis.

In the case $d = 2$, the previous definition is within a constant factor of the ordinary notion of an angle between two edges of a polygon, since only the case $\dim(F) = \dim(G) = 1$ matters when $d = 2$.

Now for the first main result of this section: we show that if there is a contractible path Π forming an angle θ, then θ^{-1} is a lower bound on the aspect ratio of at least one simplex in every possible triangulation of P.

Theorem 1 *Let F, G be two P-faces, and assume there is a contractible path from F to G. Let θ be the value of the sharpest angle between F and G. Let $\mathcal{T}$ be an arbitrary triangulation of P. Then there is a simplex T in the triangulation with a vertex lying on $F \cap G$ such that $\operatorname{asp}(T) \geq c_d/\theta$.*

Proof. Let $(\boldsymbol{x}, \boldsymbol{y}, \Pi)$ be the triple defining the sharpest angle θ between F and G, and let $\boldsymbol{z}$ be the point in $F \cap G$ to which we can contract Π. (Stating this more carefully, since sharpest angle is defined as an infimum, we should say that $(\boldsymbol{x}, \boldsymbol{y}, \Pi)$ defines an angle of size $(1+\epsilon)\theta$, where $\epsilon > 0$ is arbitrarily small. But the $1+\epsilon$ factor can be absorbed by the c_d factor.) Let H be the P-face contained in $F \cap G$ that contains $\boldsymbol{z}$. (If there is more than one P-face H satisfying $\boldsymbol{z} \in H \subset F \cap G$, choose any such H.) Let $A = \mathrm{aff}(F) \cap \mathrm{aff}(G)$. Note that $H \subset A$ since $H \subset F \cap G$.

In the triangulation of P, restrict attention to simplices of $\mathcal{T}$ that have at least one vertex on H. Call this collection of simplices $\mathcal{T}'$. Since $\mathcal{T}'$ is a finite set, there is an $\epsilon > 0$ such that every point in $\boldsymbol{v} \in P$ satisfying $\mathrm{dist}(\boldsymbol{v}, \boldsymbol{z}) \leq \epsilon$ and $\boldsymbol{v}\boldsymbol{z} \subset P$ is contained in a simplex from $\mathcal{T}'$. Then we can contract $(\boldsymbol{x}, \boldsymbol{y}, \Pi)$ toward $\boldsymbol{z}$ (i.e., replace $\boldsymbol{x}$ by $(1-\lambda)\boldsymbol{z} + \lambda\boldsymbol{x}$, $\boldsymbol{y}$ by $(1-\lambda)\boldsymbol{z} + \lambda\boldsymbol{y}$, and each point $\boldsymbol{v} \in \Pi$ by $(1-\lambda)\boldsymbol{z} + \lambda\boldsymbol{v}$ for some fixed $\lambda \in (0, 1])$ so that, without loss of generality, all of Π is covered by simplices in $\mathcal{T}'$. Note that the contraction operation does not affect the value of θ because the numerator and denominator of (2) scale by the same amount when we contract toward $\boldsymbol{z}$.

Without loss of generality, $\mathrm{dist}(\boldsymbol{x}, A) \geq \mathrm{dist}(\boldsymbol{y}, A)$; define $\alpha = \mathrm{dist}(\boldsymbol{y}, A)$. Define $\beta = \mathrm{lth}(\Pi)$. Thus, $\theta = \beta/\alpha$ is the sharpest angle. Now define a continuous piecewise linear function $f : P \to \mathbb{R}$ as follows. We first define f on the vertices of $\mathcal{T}$ as follows. For each $\mathcal{T}$-vertex $\boldsymbol{v} \in G$ we define $f(\boldsymbol{v}) = \mathrm{dist}(\boldsymbol{v}, A)$ where distance is measured in the ordinary Euclidean sense. Since $H \subset A$, this fixes $f(\boldsymbol{v}) = 0$ for vertices $\boldsymbol{v} \in H$. For all other vertices $\boldsymbol{v}$ of $\mathcal{T}$ we define $f(\boldsymbol{v}) = 0$. Notice that all vertices $\boldsymbol{v}$ of F have $f(\boldsymbol{v}) = 0$ because the intersection of F and G is contained in A. Now extend f to all of P by linearly interpolating over each simplex. This yields a uniquely determined piecewise linear function $f : P \to \mathbb{R}$. Notice that f is identically 0 on F.

Next, we claim that $f(\boldsymbol{y}) \geq \alpha$. Notice that for points on G, f is a linear interpolation of the function $\boldsymbol{u} \mapsto \mathrm{dist}(\boldsymbol{u}, A)$. This latter function is a convex function because A is convex. A linear interpolant of a convex function is always greater than or equal to the function value itself, thus $f(\boldsymbol{y}) \geq \mathrm{dist}(\boldsymbol{y}, A) = \alpha$.

On the other hand, $f(\boldsymbol{x}) = 0$ because $\boldsymbol{x} \in F$. Let $f|_\Pi$ be the restriction of f to Π; then $f|_\Pi$ is PL and continuous, and increases by α. The length of Π is β. Therefore, there is a point $\boldsymbol{u}$ where the directional derivative of f at $\boldsymbol{u}$ parallel to Π is at least α/β in magnitude. Let T be the simplex of $\mathcal{T}$ containing $\boldsymbol{u}$ (if there is more than one such T, choose arbitrarily). Note that $T \in \mathcal{T}'$ because we are assuming Π is covered by $\mathcal{T}'$. On this simplex T, since the gradient is constant, $\|\nabla f\| \geq \alpha/\beta$.

There is an analytic expression for ∇f on T as follows. Let us number the vertices of T with $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_d$ so that $\boldsymbol{v}_0$ is a vertex on H. (Recall that every simplex in $\mathcal{T}'$ has at least one vertex on H.) Let $r_i = f(\boldsymbol{v}_i)$ for $i = 0, \ldots, d$. Thus, $r_0 = 0$ because f is zero on H as noted above. Then one checks that $f(\boldsymbol{u})$ on T is given by the linear mapping $f(\boldsymbol{u}) = \boldsymbol{r}^{\mathrm{T}} M_T^{-1}(\boldsymbol{u} - \boldsymbol{v}_0)$, where M_T was defined earlier, and $\boldsymbol{r}^{\mathrm{T}}$ denotes $(r_1, \ldots, r_d)$. Then we see that ∇f on T is given by $M_T^{-\mathrm{T}}\boldsymbol{r}$, so

$$\|\nabla f\| \leq \|M_T^{-\mathrm{T}}\| \cdot \|\boldsymbol{r}\| \leq c_d \cdot (\max |r_i|)/\operatorname{minalt}(T).$$

Notice that $\max |r_i|$ is the maximum distance of a vertex of T from A, but this is at most $\operatorname{maxside}(T)$ since T has an edge from each of its vertices to $\boldsymbol{v}_0$, which lies on A. Thus,

$$\begin{aligned} \|\nabla f\| &\leq c_d \operatorname{maxside}(T)/\operatorname{minalt}(T) \\ &= c_d \operatorname{asp}(T). \end{aligned}$$

On the other hand, we showed in the previous paragraph that $\|\nabla f\| \geq \alpha/\beta$ which is the reciprocal of θ. Thus, we have proved that the aspect ratio of T is bounded below by the reciprocal of the sharpest angle between F and G. ∎

The previous result shows that the presence of a contractible path gives a useful bound that is applicable to any triangulation $\mathcal{T}$. On the other hand, the presence of an incontractible path also gives a useful bound. We start with a lemma, and then prove the main result.

Lemma 3 *Let Π be an incontractible path from $\boldsymbol{x} \in F$ to $\boldsymbol{y} \in G$. Let $\mathcal{T}$ be an arbitrary triangulation of P. Let $F_1, \ldots, F_m$ be an enumeration of all the faces (of all dimensions including d) of $\mathcal{T}$ that meet Π. Then $F_1 \cap \cdots \cap F_m = \emptyset$.*

Proof. Suppose that $F_1, \ldots, F_m$ have a common point $\boldsymbol{z}$. Since the triangulation is boundary-conforming, $\boldsymbol{z}$ lies on a common subface of F and G (because the lowest dimensional triangulation face meeting $\boldsymbol{x}$ must be a subface of F, and similarly for the lowest dimensional face meeting $\boldsymbol{y}$). Furthermore, every point $\boldsymbol{v}$ on Π is covered by a simplex that also covers $\boldsymbol{z}$. Since simplices are convex, this simplex also covers the segment $\boldsymbol{vz}$. Thus, Π is contractible to $\boldsymbol{z}$, contradicting the assumption. ∎

Theorem 2 *Let F, G be two P-faces, and let Π be an incontractible path from $\boldsymbol{x} \in F$ to $\boldsymbol{y} \in G$. Let $\mathcal{T}$ be an arbitrary triangulation of P. Then $\mathcal{T}$ contains a simplex T meeting Π such that*

$$\operatorname{minalt}(T) \leq c_d \operatorname{lth}(\Pi).$$

Proof. Let $F_1, \ldots, F_m$ be an enumeration of faces of $\mathcal{T}$ meeting Π. By the preceding lemma, $F_1 \cap \cdots \cap F_m = \emptyset$. Let the vertices of F_1 be denoted $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_s$, where $s \leq d$. Since Π meets F_1, there is a point, say $\boldsymbol{z}$, on Π that can be written as a convex combination of the vertices of F_1:

$$\boldsymbol{z} = \lambda_0 \boldsymbol{v}_0 + \cdots + \lambda_s \boldsymbol{v}_s,$$

where each λ_i is nonnegative, and $\lambda_0 + \cdots + \lambda_s = 1$. Therefore, for some i, $\lambda_i \geq 1/(s+1) \geq 1/(d+1)$. Without loss of generality, say that $\lambda_0 \geq 1/(d+1)$. Note that since the F_i's are disjoint, there is some F_i that does not contain $\boldsymbol{v}_0$. Let this other face be denoted F_2.

Let $f : P \to \mathbb{R}$ be a piecewise linear continuous function defined as follows. We set $f(\boldsymbol{v}_0) = 1$. For all other vertices $\boldsymbol{v}$ of $\mathcal{T}$, set $f(\boldsymbol{v}) = 0$. Now extend f to all of P by linear interpolation over the simplices in $\mathcal{T}$. Note that $f(\boldsymbol{z}) = \lambda_0 f(\boldsymbol{v}_0) + \cdots + \lambda_s f(\boldsymbol{v}_s)$, and hence $f(\boldsymbol{z}) \geq 1/(d+1)$. On the other hand, let $\boldsymbol{w}$ be the point where Π meets F_2; note that $f(\boldsymbol{w}) = 0$ since f is identically zero on F_2 (because f is defined to be zero on all vertices of F_2).

We now conclude the proof using the same technique as in Theorem 1. Along path Π, f is PL and continuous and decreases by at least $1/(d+1)$ (from $\boldsymbol{z}$ to $\boldsymbol{w}$). Therefore, there is a point $\boldsymbol{u}$ on Π such that f is has a directional derivative at $\boldsymbol{u}$ parallel to Π whose magnitude is at least $(1/(d+1))/\operatorname{lth}(\Pi)$. Let T be the simplex containing $\boldsymbol{u}$. Then on this simplex T, since the gradient is constant, $\|\nabla f\| \geq 1/((d+1)\operatorname{lth}(\Pi))$.

We obtain analytic expression for ∇f on T as follows. Let us number the vertices of T with $\boldsymbol{v}_0, \ldots, \boldsymbol{v}_d$. Let $r_i = f(\boldsymbol{v}_i) - f(\boldsymbol{v}_0)$ for $i = 0, \ldots, d$. Thus, $|r_i| \leq 1$ for each i. As earlier, ∇f on T is given by $M_T^{-\mathrm{T}} \boldsymbol{r}$, so

$$\|\nabla f\| \leq \|M_T^{-\mathrm{T}}\| \cdot \|\boldsymbol{r}\| \leq c_d / \operatorname{minalt}(T).$$

Combining this inequality with the inequality proved in the previous paragraph proves the theorem. ∎

10 QMG aspect ratio in terms of neighboring box sizes

In this section we begin our analysis of the aspect ratio bound for QMG. In general, we cannot establish a universal constant upper bound on the aspect ratio of the triangulation produced by QMG because if P has sharp angles, then any possible triangulation, including QMG, will have poor aspect ratio near the sharp angle as proved by Theorem 1. Thus, we want to show that the sharpest angle of any simplex generated QMG is very sharp only if the input polyhedron itself has a sharp angle.

In this section we argue that the worst-case aspect ratio produced by QMG is bounded in terms of the ratio of sizes of neighboring boxes. This section requires an understanding of the analysis in our other paper [12]. In subsequent sections, we then bound this box-size ratio in terms of the sharpest angle. From now on, we denote the sharpest angle in P by $\theta(P)$. Thus, the combination of these arguments bounds the aspect ratio of QMG in terms of $\theta(P)$.

Let B be a box. As above, we define $\mathrm{size}(B)$ to be the length of a side of B. Let B, B' be two neighboring protected boxes such that $\mathrm{co}(B)$ and $\mathrm{co}(B')$ have a common point. (Here, $\mathrm{co}(B)$ refers to the content of B at the time it became protected. The contents of two neighboring boxes might not have a common point if the boxes' common subface is completely outside P because of boundaries that cut through the boxes, or because of duplication.) Suppose $\mathrm{size}(B) \geq \mathrm{size}(B')$. These boxes have *box-size ratio* $\mathrm{size}(B)/\mathrm{size}(B')$. Let r be the maximum box-size ratio in the whole triangulation produced by QMG. We argue in this section that the worst aspect ratio in QMG is at most $c_d r$.

Consider a simplex T generated by QMG. As in [12], this simplex comes from a chain of $d+1$ nested box subfaces. Unlike [12], these box subfaces can have different sizes; in particular, the subfaces in the chain can grow in size as the dimension of the box face increases.

A consequence of the alignment condition presented in Section 6 is as follows. Let B_i, B_{i+1} be two boxes in a chain, so that $\dim(B_i) = i$ and $\dim(B_{i+1}) = i+1$. Let the close points of these boxes be $\boldsymbol{v}_i, \boldsymbol{v}_{i+1}$. Then either $\boldsymbol{v}_i = \boldsymbol{v}_{i+1}$ or else

$$\mathrm{dist}(\boldsymbol{v}_{i+1}, \mathrm{aff}(B_i)) \geq c_d \,\mathrm{size}(B_{i+1}). \tag{3}$$

The reason for this is as follows. Let C be the close face of B_{i+1}. Let B^* be the i-dimensional face of B_{i+1} that contains B_i. If C is a subface of B^*, then the alignment condition implies that B^* must also be protected, and hence $B^* = B_i$ and $\boldsymbol{v}_i = \boldsymbol{v}_{i+1}$. Else C is not a subface of B^*, which means that $\boldsymbol{v}_i$ is bounded away by $c_d \operatorname{size}(B_{i+1})$ from B^* as argued in [12], because the neighborhoods $N(\cdot)$ defined earlier create an exclusion zone around B^*.

Let M_T be the $d \times d$ matrix associated with T defined above, with columns ordered according to the chain order. Because of (3), M_T, when scaled to unit box size, satisfies analogs of the inequalities that were developed in [12] in the case of unit box size.

In particular, let S_T be the $d \times d$ diagonal matrix whose ith entry is the box side length of the ith box face in the chain defining T. Then the matrix $N = M_T S_T^{-1}$ has its columns rescaled so that each column corresponds to a difference between two vertices in a unit-size cube. Slight generalizations of the bounds proved in [12] apply to N (actually, that paper considered the transpose N^{T}). In particular, the bounds in our other paper imply that $\|N\|$ and $\|N^{-1}\|$ are at most c_d.

Since $\kappa(M_T) \leq \kappa(N)\kappa(S_T)$, we have from the last paragraph that $\kappa(M_T) \leq c_d \kappa(S_T)$. Note that all the box faces in a chain come from mutual neighboring cubes, so $\kappa(S_T) \leq r$. Therefore, $\kappa(M_T) \leq c_d r$. This argument has established the following theorem.

Theorem 3 *Let $\rho_{\mathrm{QMG}}(P)$ denote the worst-case aspect ratio produced by QMG when applied to polyhedral domain P. Then there exist two neighboring protected boxes B, B' such that $\operatorname{co}(B) \cap \operatorname{co}(B') \neq \emptyset$ and such that*

$$\rho_{\mathrm{QMG}}(P) \leq c_d \cdot \operatorname{size}(B) / \operatorname{size}(B').$$

In subsequent sections, we bound the maximum box size ratio in terms of the sharpest angle $\theta(P)$. The ultimate goal is Theorem 5 which bounds $\rho_{\mathrm{QMG}}(P)$ in terms of $\theta(P)$.

11 A bound on splitting for alignment

As we saw in the preceding section, the aspect ratio of QMG can be bounded if we can bound the number of times boxes are split. The following is the key theorem about how many times a box can be split. The proof of Theorem 4 will be the topic of this and the next few sections.

Theorem 4 *Let P be the input polyhedral region, whose sharpest angle is θ. Let B be a protected box produced by QMG. Then there exists an active box B_a that is an ancestor of B such that*

$$\mathrm{size}(B_a) \le c_d \max(1, \theta(P)^{-\phi(d)})\, \mathrm{size}(B) \tag{4}$$

where the exponent $\phi(d)$ is defined by (6) below, and such that $\mathrm{co}(B_a)$ contains an incontractible path Π satisfying $\mathrm{lth}(\Pi) \le c_d\, \mathrm{size}(B_a)$. We call B_a the **anchor** *of B.*

Recall that QMG splits boxes in both the separation and alignment stages. The purpose of this section is to show that the amount of splitting for alignment is bounded by c_d, which is one step in the proof of Theorem 4. We start with two preliminary lemmas, which lead to the main result Lemma 6 at the end of this section. That lemma is one step in the proof of Theorem 4.

Lemma 4 *Let B be a box with a neighbor B' such that $\mathrm{co}(B) \cap \mathrm{co}(B') \ne \emptyset$. There is a constant c_d such that if $\mathrm{size}(B') \le c_d\, \mathrm{size}(B)$, then $\mathrm{co}(B') \subset \mathrm{co}(B)$.*

Proof. Let $s = \mathrm{size}(B)$ and $s' = \mathrm{size}(B')$. Then $\mathrm{ex}(B)$ extends out by γs from all sides of of B. Hence $\mathrm{ex}(B)$ contains any point within ∞-norm distance γs from B. In particular, if $(1+\gamma)s' \le \gamma s$, then $\mathrm{ex}(B')$ would be completely contained in $\mathrm{ex}(B)$; let c_d in the lemma be this factor $\gamma/(1+\gamma)$, Let $\boldsymbol{x}$ be a point in $\mathrm{co}(B') \cap \mathrm{co}(B)$. Then every point in $\mathrm{co}(B')$ is reachable by a PL path in $\mathrm{co}(B')$ from $\boldsymbol{x}$. This means that all of $\mathrm{co}(B')$ is contained in the component of $P \cap \mathrm{ex}(B)$ that contains $\boldsymbol{x}$, which must be $\mathrm{co}(B)$. ∎

Lemma 5 *Let B be a box that is split for alignment: in particular, say B is split during processing of O_F in phase k for some face F. Then there exists another active box B^* created by QMG such that* (1) $\mathrm{size}(B^*) = \mathrm{size}(B)$, (2) *$B^*$ was split before the phase k alignment stage (i.e., B^* was split during phases $0, \ldots, k-1$ or in the phase k separation stage), and* (3) *there is a PL path in P from $\mathrm{co}(B)$ to $\mathrm{co}(B^*)$ of length at most $c_d\, \mathrm{size}(B)$.*

Proof. Let us first prove the lemma for the simplified version of QMG in which all the boxes are full dimensional. In this case, the alignment condition was described in Section 6 as follows: B is split for alignment because its high-priority close face, say C, is not completely covered by (full-dimensional) boxes in the extended orbit of F the same size or larger when B is processed.

Consider the collection of boxes obtained by taking all boxes produced by any step of simplified QMG that are same size as B. Consider also all those boxes larger than B that are leaf boxes (i.e., protected). Notice that this collection of boxes, say Q, completely covers the input domain; there could some parts of $\mathbb{R}^d$ double-covered because of duplication.

Let the enumeration of all boxes in Q that cover C be denoted $B_1, \ldots, B_m$: exclude B itself from this enumeration. Since F is close to C, F meets $N(C)$ and hence also $\mathrm{co}(B)$. Let $\boldsymbol{x}$ be a point where F meets $N(C)$. Among $B_1, \ldots, B_m$, consider only those B_i such that $\boldsymbol{x} \in \mathrm{co}(B_i)$. The case when $\boldsymbol{x} \notin \mathrm{co}(B_i)$ could only occur because of duplication; there could be duplicates of neighbors of B that do not contain $\boldsymbol{x}$.

Rename the remaining boxes again as $B_1, \ldots, B_m$; note that these boxes together with B must cover C. Since the alignment condition does not hold for B, one of them, say B_i, is already split at the time B is processed. The box that is already split must have the same size as B (i.e., it cannot be one of the larger boxes in Q, because those boxes are all protected). Without loss of generality, B_i is the earliest box among $B_1, \ldots, B_m$ to be split.

Case 1 is that this box B_i was split during a phase $0, \ldots, k-1$, or during phase k separation. In this case the lemma is proved with $B^* = B_i$; note that the two boxes contain $\boldsymbol{x}$ in their content so that condition (3) of the lemma is trivial.

The remaining case, Case 2, is that B_i is split during the phase k alignment stage before B is processed. Rename B_i as B'. Note that since $\mathrm{co}(B')$ meets F, then B' must be in O_F. Since B' is the first box among $B_1, \ldots, B_m$ to be split during phase k alignment, C is still covered by boxes in the orbit of the same size or larger at the time the alignment condition is checked for B'. Let C' be the close face of B'. Note that C' is not covered by boxes in the orbit of the same size or larger, since we are assuming B' is split for alignment in phase k. It is not possible that $\dim(C') \geq \dim(C)$ because then C would have higher priority than C' and hence C' would not be selected as the close face of B'. (Recall that the priority rule favors faces of lower dimension, and among faces of the same dimension, favors faces completely covered by boxes in the orbit.) Thus, $\dim(C') < \dim(C)$. Start the proof of this lemma over again with B' and C'. In other words, consider the boxes in the quadtree at the level of B' that cover C'. Either for B' we will "exit" this argument in Case 1 (i.e., we find a box B^* that satisfies the lemma for B'), or we will have to restart the argument another time.

But note that each time we restart the above argument, the dimension of the close face in question decreases by 1. Thus, we can repeat the argument at most d times before terminating at Case 1. Let the sequence of boxes constructed by repeating this argument be $B, B', B'', \ldots, B^{(r)}, B^*$, where B^* is a box split before phase k alignment. Note that $r \le d$ as just mentioned. Also, all boxes in this sequence are the same size, and $B^{(i)}$ is adjacent to $B^{(i+1)}$ for each i. Furthermore, $\mathrm{co}(B^{(r)})$ has a common point with $\mathrm{co}(B^*)$, and all of $B^{(1)}, \ldots, B^{(r)}$ are in O_F. This means that we can find a PL path in P from B to $B^{(r)}$ by traversing F through each box. (Recall that a box in O_F cannot meet any boundaries of F in its content.) The length of the PL path constructed in this manner is at most $c_d\,\mathrm{size}(B)$. This proves the lemma.

If we wanted to extend this proof to the case of the complete version of QMG (including lower dimensional boxes), then we would use the same proof as above, except that we have to restate the meaning of "completely covered" in terms of the weight system mentioned in Section 7. Because we have incompletely described the weight system and skipped the lemmas showing that lower-dimensional boxes do indeed act like proxies for the full-dimensional boxes that contain them, we do not have enough machinery to prove this lemma in the general case so we merely assert it. ∎

The preceding lemma now leads to the main result of this section, which says that during splitting for alignment, boxes can become only a constant factor smaller.

Lemma 6 *Let B be a box that results from splitting for alignment during the processing of O_F for some k-face F. Then B is descended from an active box B' at the start of the phase k alignment stage such that* $\mathrm{size}(B') \le c_d\,\mathrm{size}(B)$.

Proof. Let B_0 be a parent of B, so that B_0 was split for alignment during the processing of O_F and so that $\mathrm{size}(B_0) = 2\,\mathrm{size}(B)$. By the preceding lemma, there is another box B^* that was either protected from an earlier phase or was split for separation, such that there is a PL path Π in P from $\mathrm{co}(B^*)$ to $\mathrm{co}(B_0)$ of length at most $c_d\,\mathrm{size}(B_0)$.

Let B' be the ancestor of B_0 at the beginning of phase k alignment in O_F. Observe that there is a constant χ_d depending on d such that if B' satisfied $\text{size}(B') \geq \chi_d \, \text{size}(B_0)$, then $\text{co}(B')$ would contain $\text{co}(B^*)$ as a subset. This follows from the same proof technique used for Lemma 4; in particular, if B' were sufficiently larger than B_0, it would contain the whole path Π and also $\text{co}(B^*)$.

On the other hand, it is impossible that $\text{co}(B')$ contains $\text{co}(B^*)$. This is because B^* was split for separation in phase k or was split in a phase earlier than k. Whatever P-faces caused B^* to be split would cause B' to be crowded, and hence B' could not end up in O_F.

Thus, we conclude $\text{size}(B') < \chi_d \, \text{size}(B_0)$, which proves the lemma. ∎

This lemma shows that all splitting for alignment can be lumped into the factor c_d in (4).

12 Splitting boxes for weak crowding

Recall that a box is split for separation if and only if it is crowded. Recall also that there are two ways that B can be crowded in phase k: (1) $\text{co}(B)$ contains a P-face of dimension $k-1$ or lower, or (2) $\text{co}(B)$ contains a P-face F of dimension k, and another P-face G of dimension k or greater that is not a superface of F.

We call the former "weak crowding" and the latter "strong crowding." In this section we show that all splitting for weak crowding can also be lumped into the factor c_d in (4), which is another step toward proving Theorem 4.

Lemma 7 *Let B be a box that is split for weak crowding, that is, in the phase k separation stage,* $\text{co}(B)$ *meets a P-face of dimension $k-1$ or lower. Then B has an ancestor B_0 that is an active box at the beginning of phase k such that* $\text{size}(B) \geq c_d \, \text{size}(B_0)$.

Proof. The assumption implies that the P-face F of dimension $k-1$ or lower passes through $\text{ex}(B)$. Let B_0 be the ancestor of B from the beginning of phase k. We claim that B_0 can be at most a constant factor c_d larger than B. This is because the expansion factors for $\text{ex}(B)$ and $\overline{\text{ex}}(B)$ are off by a constant c_d. Recall that $\overline{\text{ex}}(B)$ was defined in Section 6 and is applied here with respect to face F. Therefore, the ancestor of B_0, if it is much larger than B, would contain this P-face F in $\overline{\text{ex}}(B_0)$. Recall that if F meets $\overline{\text{ex}}(B_0)$, then B_0 has a subface close to F. Hence B_0 would have been protected in phase $\dim(F)$ which is less than k, or would have been split for separation. Thus, B_0 is at most c_d larger than B. ∎

Thus, all splitting for weak-crowding can also be lumped into the factor c_d in (4).

13 Splitting for strong crowding

In this section we analyze splitting for strong crowding and finally prove Theorem 4. Recall that "strongly crowded" means that there is a P-face F of dimension k in $\mathrm{co}(B)$, and another P-face G of dimension $l \geq k$ in $\mathrm{co}(B)$ that is not a superface of F. From now on, we say that G is "foreign" to F if G is not a superface of F. We say that two points $\boldsymbol{x}$ and $\boldsymbol{y}$ are "visible" to each other with respect to P if segment $\boldsymbol{xy}$ lies in P. We start with two lemmas about visibility.

Lemma 8 *Let P be a k-dimensional polyhedral domain in $\mathbb{R}^d$, and let C be a convex subset of $\mathbb{R}^d$. Suppose $P \cap C$ is not empty, and let U be a component of $P \cap C$. Suppose that U meets a P-face F, and that U does not meet any faces of P that are foreign to F. Then every point in U is visible to every point in $F \cap U$, where "visibility" is with respect to U.*

See Fig. 9 for an illustration of this lemma.

Proof. Let $\boldsymbol{x}$ be a point in $F \cap U$ and $\boldsymbol{y}$ a point in U. Consider the segment $L = \boldsymbol{xy}$; suppose that this segment is not contained in U. We will derive a contradiction. Since $\boldsymbol{y} \in U$ and U is closed, there must some point $\boldsymbol{z} \in L$ different from $\boldsymbol{x}$ such that $\boldsymbol{z}$ is in U, but there is a sequence of points $\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots$ lying on L and converging to $\boldsymbol{z}$ that are not in U. Note that all of these points lie in C because C is convex and L joins two points in C. Thus, since these points are not in $P \cap C$, we conclude that they are not in P. This means that there is at least one facet H of P (where "facet" refers to a face of dimension $k-1$) passing through $\boldsymbol{z}$ such that $\mathrm{aff}(H)$ does not contain L as a subset.

But this is impossible, because every facet of P meeting U in this component is a superface of F by assumption. This means in particular that for H in the last paragraph, $\boldsymbol{x} \in \mathrm{aff}(H)$. But since $\mathrm{aff}(H)$ is convex and contains $\boldsymbol{x}$ and $\boldsymbol{z}$, it also contains L. ∎

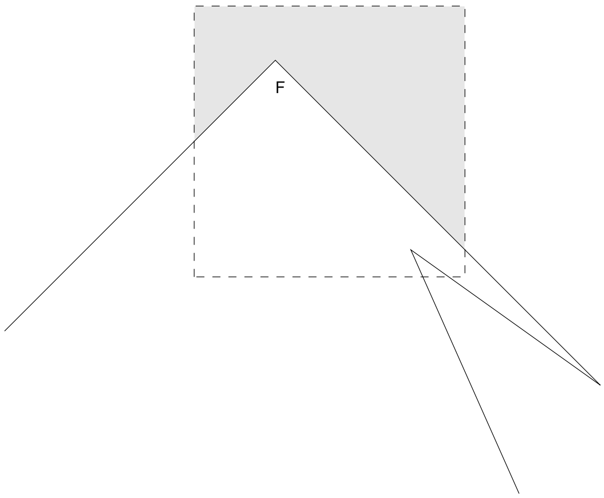

Figure 9: This figure illustrates Lemma 8 in the case $k = d = 2$. Face F is a single vertex. The boundary of P is the solid line. The convex set C is the dashed square in the figure. The shaded region is U. Notice that every point in U is visible to F.

Lemma 9 *Let P be a k-dimensional polyhedral domain in $\mathbb{R}^d$, and let C be a convex subset of $\mathbb{R}^d$. Suppose $P \cap C$ is not empty, and let U be a component of $P \cap C$. Suppose that U meets a P-face F, and suppose that U also meets a face G of P foreign to F. Then for any point $\boldsymbol{x} \in F \cap U$, there is a point $\boldsymbol{y} \in U$ that is visible to $\boldsymbol{x}$ (with respect to U) such that $\boldsymbol{y}$ lies on a P-face foreign to F (which may or may not be G).*

See Fig. 10 for an illustration.

Proof. For $\lambda \in [0, 1]$, let $C(\lambda)$ denote the contraction by λ of C toward $\boldsymbol{x}$ (i.e., $\boldsymbol{v} \in C$ iff $\lambda\boldsymbol{v} + (1 - \lambda)\boldsymbol{x} \in C(\lambda)$). Let $U(\lambda)$ be the component of $C(\lambda) \cap P$ that contains $\boldsymbol{x}$. Find the parameter value $\lambda^* > 0$, such that $U(\lambda^*)$ still meets a foreign face, but $U(\lambda^* - \epsilon)$ meets no faces for to F for all small $\epsilon > 0$. By the preceding lemma, every point in $U(\lambda^* - \epsilon)$ is visible to $\boldsymbol{x}$. Since the set of points visible to $\boldsymbol{x}$ is closed, this means that every point in $U(\lambda^*)$ is also visible to $\boldsymbol{x}$, and this set includes a point from a foreign face. ∎

We now conclude the proof of Theorem 4. Let B be a protected box that is produced by QMG. Write down its sequence of ancestors $B_0, B_1, \ldots, B_r$

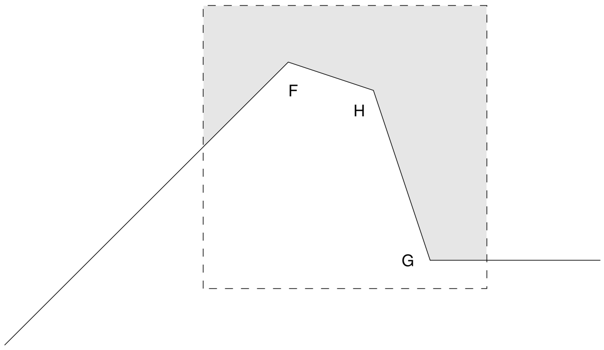

Figure 10: This figure illustrates Lemma 9 for $d = 2$. Face F is a single vertex. The convex set C is the dashed square in the figure. The shaded region is U. Note that there is a vertex G of P that is foreign to F. This means that there is a face, namely vertex H in U, that is also foreign to F but is visible to $F \cap U$.

where $B_r = B$ and B_0 is the top box. This sequence is not necessarily unique if B_r is not full dimensional, in which case any sequence of ancestors will do. Now, delete boxes B_i in this sequence such that B_{i+1} arises from B_i via subface launching. Denote the new list $B_0, \ldots, B_r$ again. Each box is now a factor 2 smaller than its predecessor. From this list, delete boxes that are split either for alignment or for weak crowding, and denote the new list again as $B_0, B_1, \ldots, B_r$. This new list contains boxes that are split only for strong crowding, as well as the protected box B_r which is not split. By Lemma 6 and Lemma 7, in this new sequence of boxes, each box differs from its predecessor in size by a factor at most c_d.

Each box $B_0, \ldots, B_{r-1}$ is split during some phase 0 to $d-1$. (There cannot be any strong crowding in phase d by definition of strong crowding.) Therefore, mark the location where each phase begins and ends in the sequence $B_0, \ldots, B_{r-1}$. This divides the sequence $B_0, \ldots, B_{r-1}$ into "periods," where the kth period consists of boxes split during phase k.

Now we will subdivide each period into *subperiods*, using the following procedure. Focus on one particular period k, and suppose it starts at B_l and ends with B_m (i.e., B_l is the first box of the sequence split in phase k, and

B_m is the last). Since B_m is strongly crowded, there is a k-face of P, say F, meeting $\mathrm{co}(B_m)$, and there is another face G foreign to F meeting $\mathrm{co}(B_m)$. Pick a point $\boldsymbol{x} \in F \cap \mathrm{co}(B_m)$. Without loss of generality, by Lemma 9, we can assume that there is a segment in $\mathrm{co}(B_m)$ from $\boldsymbol{x}$ to a point $\boldsymbol{y} \in G$ (else choose a different G). Similarly, without loss of generality, $\boldsymbol{y}$ is not in any proper subface of G. (If $\boldsymbol{y}$ is in a proper subface of G, simply reselect G to be the subface: because G is foreign to F, every subface of G is also foreign to F).

We will construct a PL path Π in $\mathrm{co}(B_m)$. The first segment of the path is $\boldsymbol{xy}$. Consider whether G has any boundary faces that meet $\mathrm{co}(B_m)$. If not, then the construction of Π is complete, i.e., we let $\Pi = \boldsymbol{xy}$. The other case is that a boundary of G meets $\mathrm{co}(B_m)$. By Lemma 9, there is a segment in $\mathrm{co}(B_m)$ from $\boldsymbol{y}$ to a boundary face of G. (In this application of Lemma 9, the "polyhedral domain" in the lemma is G itself. Note that a boundary of G is foreign to G, i.e., it is not a superface.) Append this new segment to Π. We can continue in this manner until we reach a point which we will now be reassigned the name $\boldsymbol{y}$, such that $\boldsymbol{y}$ lies on P-face foreign to F, which we rename G, such that G does not have any boundaries that meet $\mathrm{co}(B_m)$. Note that $\mathrm{lth}(\Pi)$ is at most $d \cdot \sqrt{d}(1+\gamma)\,\mathrm{size}(B_m)$. The factor $\sqrt{d}(1+\gamma)\,\mathrm{size}(B_m)$ is the diameter of $\mathrm{ex}(B_m)$ (in the worst case, when B_m is full-dimensional) and hence is the maximum length of any segment in $\mathrm{co}(B_m)$, and the factor d comes from the fact that Π has at most d segments in it. This is because in the preceding construction of Π, each time a new segment is added, the dimension of the boundary face in question decreases by at least 1. Thus, $\mathrm{lth}(\Pi) \le c_d\,\mathrm{size}(B_m)$.

Note that $\Pi \subset \mathrm{co}(B_i)$ for each $i = l, \ldots, m$ since $\mathrm{co}(B_m) \subset \mathrm{co}(B_{m-1}) \subset \cdots \subset \mathrm{co}(B_l)$. On the other hand, in an ancestor of B_m, it might be possible to extend Π with one or more additional segments so that it reaches a lower-dimensional boundary face of G. Find the lowest numbered box B_q (largest in size) in the period $B_l, \ldots, B_m$ such that it is *not* possible to extend Π to a P-face of lower dimension that is a boundary of G. We will say that the subsequence $B_q, B_{q+1}, \ldots, B_m$ is one **subperiod** of period k. Each box in this subperiod is associated with the quintuple $(\boldsymbol{x}, \boldsymbol{y}, F, G, \Pi)$ defined in the last paragraph. If $q = l$, then we are done; this is the only subperiod of period k.

Else suppose not; suppose $q > l$. Then in B_{q-1} it is possible to extend Π to reach a face of lower dimension than what was reached in B_q. Extend Π with one or more additional segments, yielding a path Π' to a face G'

that is a proper subface of G. Now we repeat the above argument to find the predecessor of B_{q-1}, say $B_{q'}$ such that Π cannot be extended in $B_{q'}$ but can be extended in $B_{q'-1}$ (or else $q' = l$), and we let $B_{q'}, B_{q'+1}, \ldots, B_{q-1}$ be another subperiod of period k, associated with $(\boldsymbol{x}, \boldsymbol{y}', F, G', \Pi')$. Continue in this manner until we finally get back to B_l. The maximum number of subperiods in period k is seen to be $d-k$. The reason is that each time we back up to a new subperiod, the dimension of the face reached by Π decreases by at least 1. The maximum possible dimension of G initially is $d-1$, and the minimum possible dimension is k (because no box among $B_l, \ldots, B_m$ is weakly crowded).

Thus, we have divided the sequence of boxes $B_0, \ldots, B_{r-1}$ into at most d periods numbered $0, \ldots, d-1$, and period k is divided into at most $d-k$ subperiods. We have also associated with each box a choice of $(\boldsymbol{x}, \boldsymbol{y}, F, G, \Pi)$.

Now we can classify each box in the sequence according to whether its associated path Π is contractible or not. Let B_a be the highest numbered (smallest) box in $B_0, \ldots, B_{r-1}$ such that its associated path is incontractible. (Notice that B_a exists because B_0 certainly satisfies this condition.) Thus, B_a satisfies the conditions of Theorem 4 that it contains an incontractible path of length at most $c_d\,\mathrm{size}(B_a)$, and that it is an ancestor of B. All that remains is to establish (4). Note that the anchor box must be the last one in its subperiod, because all boxes in a subperiod have the same associated path. From now on, we consider only the portion $B_a, \ldots, B_r$ of the original sequence and forget about $B_0, \ldots, B_{a-1}$.

We start with the following intermediate result. Let B_m be a box in the sequence with $m > a$, and suppose it is contained in a subperiod beginning with B_q (so $a < q \leq m$). Then we claim

$$\mathrm{size}(B_m) \geq c_d \cdot \min(\theta(P), 1) \cdot \mathrm{size}(B_q). \tag{5}$$

Let us assume that B_m is a proper descendent of B_q because when $m = q$, (5) is trivially true. Let $(\boldsymbol{x}, \boldsymbol{y}, F, G, \Pi)$ be the quintuple associated with B_m; by definition of subperiod, the same quintuple is associated with B_q. Since Π is contractible (by choice of B_a), $F \cap G$ is nonempty. Let A denote $\mathrm{aff}(F) \cap \mathrm{aff}(G)$. We claim that A does not meet $\mathrm{ex}(B_q)$. See Fig. 11 for an illustration of the items constructed in the proof of this claim. Note that no boundary face of F meets $\mathrm{co}(B_q)$, because $\mathrm{co}(B_q)$ does not meet any P-faces of dimension $k-1$ or less (because it is not weakly crowded.) Thus, $\mathrm{co}(B_q)$ must contain all of $\mathrm{aff}(F) \cap \mathrm{ex}(B_q)$. Similarly, $\mathrm{co}(B_q)$ does not contain any

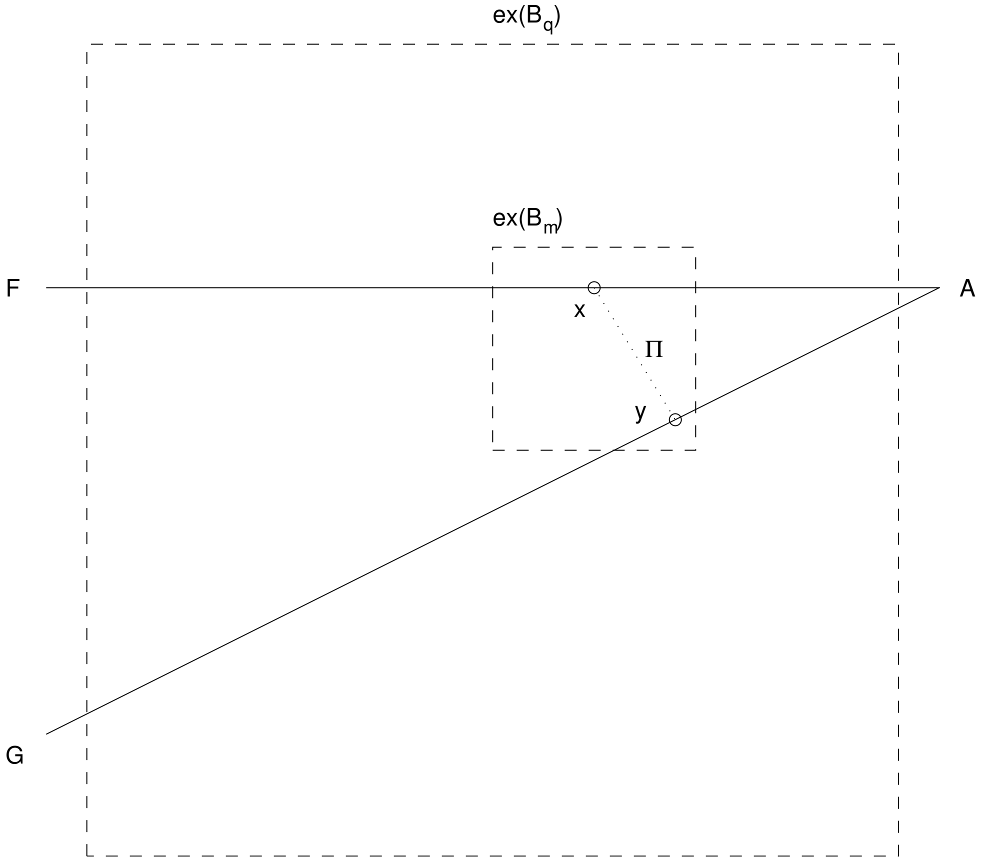

Figure 11: This is an illustration of the items $\boldsymbol{x}, \boldsymbol{y}, F, G, \Pi, A$ arising in the proof of (5). The boundaries of $\mathrm{ex}(B_q)$ and $\mathrm{ex}(B_m)$ are the dashed lines. In this figure A is zero-dimensional and is equal to $F \cap G$, but in general A will be a superset of $F \cap G$.

boundary faces of G by construction of Π (else Π could be extended). Thus, $\mathrm{co}(B_q)$ also contains all of $\mathrm{aff}(G) \cap \mathrm{ex}(B_q)$. Suppose that $\mathrm{aff}(F) \cap \mathrm{aff}(G)$ met $\mathrm{ex}(B_q)$; then $\mathrm{co}(B_q)$ would have to contain all of $\mathrm{aff}(F) \cap \mathrm{aff}(G) \cap \mathrm{ex}(B_q) = A \cap \mathrm{ex}(B_q)$ by the foregoing argument. In particular, F and G would meet in $\mathrm{co}(B)$ at all points in $A \cap \mathrm{ex}(B_q)$. But this is impossible, because neither has any boundaries in $\mathrm{co}(B_q)$.

Since A does not pass through $\mathrm{ex}(B_q)$, there is a lower bound of the form $c_d \,\mathrm{size}(B_q)$ on the distance from A to $\mathrm{ex}(B_m)$. This is because $\mathrm{ex}(B_q)$ extends a small fraction c_d multiplied by $\mathrm{size}(B_q)$ beyond $\mathrm{ex}(B')$ for any proper descendant B' of B_q. In particular, this means $\mathrm{dist}(\boldsymbol{x}, A) \geq c_d \,\mathrm{size}(B_q)$ and $\mathrm{dist}(\boldsymbol{y}, A) \geq c_d \,\mathrm{size}(B_q)$. On the other hand, $\mathrm{lth}(\Pi) \leq c_d \,\mathrm{size}(B_m)$ as argued above. Thus, in the definition of sharp angle (2) applied to $(\boldsymbol{x}, \boldsymbol{y}, \Pi)$, we see that F and G make an angle less than or equal to $c_d\,\mathrm{size}(B_m)/\,\mathrm{size}(B_q)$. Since $\theta(P)$ is the sharpest angle,

$$\theta(P) \le c_d\,\mathrm{size}(B_m)/\,\mathrm{size}(B_q).$$

This equation proves (5).

We now deduce (4) from (5). If B_q is the beginning of a subperiod with $q > a$, and B_m is its end, then $\mathrm{size}(B_q) \le c_d \max(1, \theta(P)^{-1})\,\mathrm{size}(B_m)$ from (5). Thus, if ϕ stands for the total number of subperiods between B_{a+1} and B_r, then $\mathrm{size}(B_a) \le (c_d \max(1, \theta(P)))^{-\phi}\,\mathrm{size}(B_r)$, where $c_d^{-\phi}$ accounts for the factor in box size shrinkage between the end of one subperiod and the beginning of the next, and $\theta(P)^{-\phi}$ accounts for box shrinkage within the ϕ subperiods. Since there are at most $d-k$ subperiods in period k, the total number of subperiods ϕ can be bounded $\phi(d) = d(d+1)/2$. Thus, $\mathrm{size}(B_a) \le c_d \max(1, \theta(P)^{-d(d+1)/2})\,\mathrm{size}(B_r)$ (where we have renamed $c_d^{-\phi(d)}$ as c_d).

In fact, we can immediately improve this estimate on $\phi(d)$ with the following observation. Note that if B_i is in period 0, then its path Π constructed above cannot be contractible. This is because F in phase 0 is a vertex and hence is disjoint from any foreign face G. Thus, the anchor box B_a is either the last box of period 0 or is in a later period. This means that the only subperiods that matter are in period 1 or later. Thus, we can improve the estimate to $\phi(d) = (d-1)d/2$.

We can further improve the estimate to (6) below with the following more complicated analysis. We claim that in period 1, a single subperiod suffices. The proof is as follows. Assume that the anchor box is in period 0 or 1 (else we would not need to include subperiods of period 1 in the count of ϕ, so (6) holds already). Let us review why we constructed subperiods in the first place. Let B_m be the box at the end of a subperiod, let $(\boldsymbol{x}, \boldsymbol{y}, F, G, \Pi)$ be its associated quintuple such that Π is contractible, and let B_q be the first box in the subperiod ending at B_m. In the above derivation of (5), we used the fact $\mathrm{aff}(F) \cap \mathrm{aff}(G)$ cannot pass through $\mathrm{ex}(B_q)$. To derive this fact, we needed to know that G does not have any boundaries in $\mathrm{co}(B_q)$. The above method of constructing subperiods indeed assures that G does not have boundaries in $\mathrm{co}(B_q)$.

But consider the special case of period 1, so that $\dim(\mathrm{aff}(F)) = 1$. Let B_m be the last box in period 1, and redefine B_q to be the first box of period 1, or the child of the anchor B_a, whichever comes later. Let $(\boldsymbol{x}, \boldsymbol{y}, F, G, \Pi)$

be the quintuple for B_m. Since $\dim(F)=1$, $\dim(\mathrm{aff}(F)\cap\mathrm{aff}(G))$ is either 0 or 1. The case when $\dim(\mathrm{aff}(F)\cap\mathrm{aff}(G))=1$ cannot occur by the way we construct Π. In particular, if $\dim(\mathrm{aff}(F)\cap\mathrm{aff}(G))=1$ then $\mathrm{aff}(F)\subset\mathrm{aff}(G)$, which means that either G is a superface of F (contradicting the choice of G as foreign face) or that G has a boundary in $\mathrm{co}(B_m)$ (contradicting the fact that Π reaches a face of minimal dimension).

The other case is that $\dim(\mathrm{aff}(F)\cap\mathrm{aff}(G))=0$, i.e., $\mathrm{aff}(F)\cap\mathrm{aff}(G)$ is a single point $\{\boldsymbol{v}\}$. Since Π is contractible, F and G have a common subface which must therefore be $\{\boldsymbol{v}\}$ itself. Thus $\{\boldsymbol{v}\}$ is a face of P. But since B_q is not weakly crowded, $\mathrm{co}(B_q)$ cannot contain a 0-dimensional face of P. Thus, without making any assumption about whether G has boundaries in $\mathrm{co}(B_q)$, we have determined that $\mathrm{aff}(F)\cap\mathrm{aff}(G)$ does not pass through $\mathrm{ex}(B_q)$. Therefore, a single subperiod suffices for period 1.

Thus, finally, we have the following improved estimate for ϕ, which is the total number of subperiods after B_a:

$$\phi(d)=(d-1)(d-2)/2+1. \tag{6}$$

This concludes the proof of Theorem 4.

Notice that by combining Theorem 2 and Theorem 4 we immediate obtain the following corollary.

Corollary 1 *Let B be a protected box generated by QMG. Then there exists an ancestor B_a of B such that* (4) *holds and such that for any triangulation $\mathcal{T}$ of P, there is a simplex T meeting* $\mathrm{co}(B_a)$ *such that* $\mathrm{minalt}(T)\le c_d\,\mathrm{size}(B_a)$.

In fact we can strengthen Corollary 1 (though not Theorem 4) in the case $d=2$. The strengthened version of Corollary 1 asserts that

$$\mathrm{size}(B_a)\le c\,\mathrm{size}(B), \tag{7}$$

holds when $d=2$, in place of (4) (i.e., the factor of $\theta(P)^{-1}$ goes away). The argument for this strengthening is as follows. In the following argument, c denotes an absolute constant whose value may change from formula to formula.

Consider how the factor $\theta(P)^{-1}$ arises in the first place. Let B be a protected box and B_a its anchor. This factor comes when anchor B_a is from period 0, and then in period 1 we split a strongly crowded box that has a

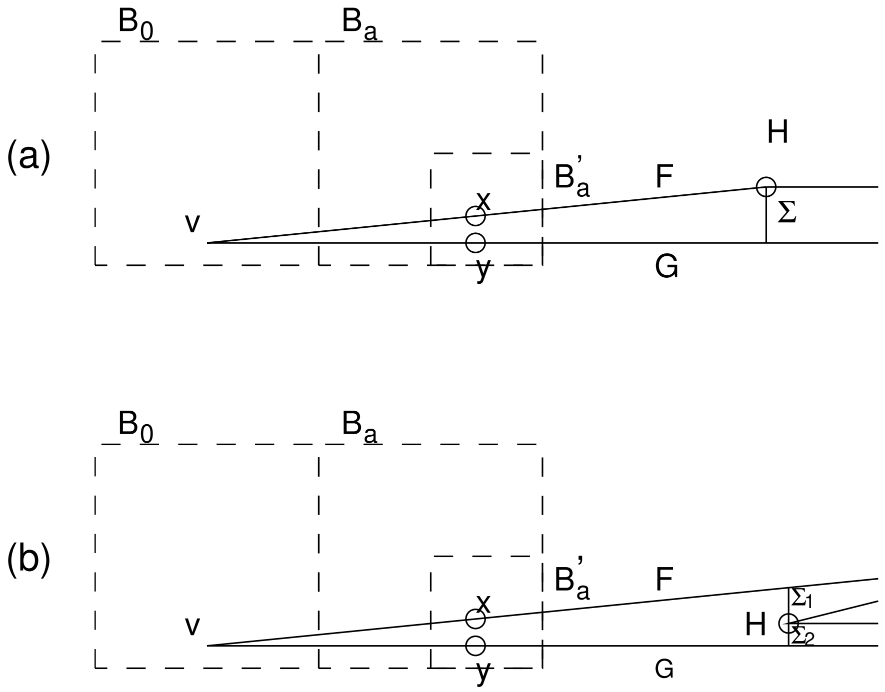

Figure 12: In this figure, we show that faces arising in the proof of the strengthened version of Corollary 1 when $d = 2$. Box B and path Π are not depicted; both are enclosed in B'_a in the figure.

contractible path that defines a sharp angle θ. Since (4) and (7) are equivalent when $\theta(P)$ is large, let us assume that $\theta \leq 0.1$. Let B'_a to be the last box that is split for strong crowding at the end of period 1. Clearly (7) holds for this choice of B'_a (since splitting for alignment, as well as all splitting in phase 2, incurs only an additional factor c). The contractible path Π in B'_a is from $\boldsymbol{x}$ to $\boldsymbol{y}$. See Fig. 12. We now must prove that for an arbitrary triangulation $\mathcal{T}$, there is a simplex T meeting Π satisfying $\operatorname{minalt}(T) \leq c_d \operatorname{size}(B'_a)$.

In the figure, B_0 denotes a protected box for $\boldsymbol{v}$, the common subface of F and G. Note that there must be an incontractible path from $\boldsymbol{v}$ to a foreign face H whose length at most a factor c more than $\operatorname{size}(B_0)$ by the preceding analysis. Thus, $\boldsymbol{x}$ and $\boldsymbol{y}$ must both be separated from $\boldsymbol{v}$ by at least $c_d \operatorname{dist}(\boldsymbol{v}, H)$. There are two possible ways to choose H; either it is a subface of one of F or G, which is (a) in the figure, or it is does not meet F and G, which is (b) in the figure.

Let $\mathcal{T}$ be an arbitrary triangulation. Let $T_1, \ldots, T_s$ be the triangles of T that meet Π. Take two cases. In the first case, suppose that at least one of $T_1, \ldots, T_s$, say T_1, meets the segment denoted by Σ in (a) or that it meets the PL path $\Sigma_1 \cup \Sigma_2$ in (b). Since $\Sigma, \Sigma_1, \Sigma_2$ are all incontractible, this means by Theorem 2 that $\mathrm{minalt}(T_1) \leq c\,\mathrm{lth}(\Sigma)$ for (a) or $\mathrm{minalt}(T_1) \leq c\,\mathrm{lth}(\Sigma_1 \cup \Sigma_2)$ for (b). But now it is clear that $\mathrm{lth}(\Sigma), \mathrm{lth}(\Sigma_1), \mathrm{lth}(\Sigma_2)$ are all at most $c\theta\,\mathrm{dist}(\boldsymbol{v}, H)$, that is, less than or equal to $c\,\mathrm{size}(B'_a)$. Thus, $\mathrm{minalt}(T_1) \leq c\,\mathrm{size}(B'_a)$, proving the corollary.

The other case is that none of $T_1, \ldots, T_s$ meet Σ in (a) or $\Sigma_1 \cup \Sigma_2$ in (b). But this means all of $T_1, \ldots, T_s$ are contained in the polygonal region bounded by $F \cup G \cup \Sigma$ in (a) or $F \cup G \cup \Sigma_1 \cup \Sigma_2$ in (b). The width of this polygonal region is at most $c\,\mathrm{size}(B'_a)$, so any triangle in this region has minalt at most $c\,\mathrm{size}(B'_a)$. This concludes the proof that Corollary 1 may be strengthened in the case $d = 2$.

14 Maximizing the minimum altitude

In this section we consider the problem of computing a triangulation that maximizes the minimum altitude. Although this is not the problem for which QMG is intended, we nonetheless can obtain an interesting consequence from Corollary 1. Suppose we want to compute the triangulation of P that maximizes the minimum altitude. In other words, for a triangulation $\mathcal{T}$ of P, define

$$\mu_{\mathcal{T}}(P) = \min\{\mathrm{minalt}(T) : T \in \mathcal{T}\}$$

and then consider the triangulation $\mathcal{T}^*$ that solves

$$\mu(P) = \max\{\mu_{\mathcal{T}}(P) : \mathcal{T} \text{ is a triangulation of } P\}.$$

It follows from Corollary 1 that QMG solves this problem to within a factor $c_d\theta(P)^{-\phi(d)}$ for $d > 2$ and within a factor c (a universal constant) when $d = 2$. This is because the minimum altitude among all triangles produced by QMG is within a factor c_d of the smallest protected box generated by QMG. But the smallest protected box is within a factor of $c_d\theta(P)^{-\phi(d)}$ of the minimum altitude of any possible triangulation by the corollary.

Thus, in the case $d = 2$, $\mu_{\mathrm{QMG}}(P) \geq c\mu(P)$, and for $d = 3$, $\mu_{\mathrm{QMG}}(P) \geq c\min(1, \theta(P)^2)\mu(P)$. In the case $d = 2$, more is known about this problem. In particular, the algorithm of Bern, Dobkin and Eppstein [2] produces a triangulation $\mathcal{T}$ also satisfying $\mu_{\mathcal{T}}(P) \geq c\mu(P)$, and also $\mathcal{T}$ has an optimal (linear) number of triangles.

In the case $d = 3$, much less is know. For instance, we do not know whether the bound $c\theta(P)^{-2}$ is tight for QMG. We have constructed an example where the minimum altitude of the triangulation produced by QMG is off from $\mu(P)$ by a factor $c\theta(P)^{-1}$, but we have not found an example attaining the bound $c\theta(P)^{-2}$.

Another open question concerns a geometric characterization of $\mu(P)$. It follows from the optimality of QMG in case $d = 2$ that $\mu(P)$ is within a constant factor of the minimum geodesic distance between P-faces that do not meet each other. Is there a similar simple geometric characterization of $\mu(P)$ for $d = 3$? For any d, Theorem 2 implies that the minimum geodesic distance between two non-meeting faces of P is an upper bound on $\mu(P)$ (to within a constant c_d), but it is not known whether this bound is tight.

15 A bound on the QMG aspect ratio

This section establishes the optimality of the QMG aspect ratio using Theorem 4 and Corollary 1.

Theorem 5 *Let $\rho_{\mathrm{QMG}}(P)$ denote the worst-case aspect ratio produced by QMG when applied to polyhedral domain P. Then*

$$\rho_{\mathrm{QMG}}(P) \leq c_d \max(1, \theta(P)^{-\phi(d)}).$$

Proof. Let B_1, B_2 be two neighboring protected boxes such that $\mathrm{co}(B_1)$ and $\mathrm{co}(B_2)$ have a common point. Assume that B_1 is protected in phase k and lies in O_F, and B_2 is protected in phase l with $l \geq k$ and lies in O_G, and assume $\mathrm{size}(B_1) > \mathrm{size}(B_2)$. By Theorem 3, it suffices to obtain an upper bound on $\mathrm{size}(B_1)/\mathrm{size}(B_2)$.

We can assume that $\mathrm{co}(B_2) \subset \mathrm{co}(B_1)$. If this relation did not hold, then by Lemma 4 we could immediately conclude that there is a bound of the form c_d on $\mathrm{size}(B_1)/\mathrm{size}(B_2)$.

By Theorem 4, B_2 has an anchor B_a containing an incontractible path such that

$$c_d \max(\theta(P)^{-\phi(d)}, 1)\,\mathrm{size}(B_2) \geq \mathrm{size}(B_a). \tag{8}$$

Because B_1 is protected, every point in $\mathrm{co}(B_1)$ is visible to every point in $F \cap \mathrm{co}(B_1)$ by Lemma 8, and this includes $\mathrm{co}(B_2)$ as well. Therefore, any PL path inside $\mathrm{co}(B_1)$ is contractible to any point of $F \cap \mathrm{co}(B_1)$. Since $\mathrm{co}(B_a)$ contains an incontractible path, it cannot be a a subset of $\mathrm{co}(B_1)$. This means that $\mathrm{size}(B_a) \geq c_d\,\mathrm{size}(B_1)$. Combining this with (8) shows that $\mathrm{size}(B_1)/\mathrm{size}(B_2) \leq c_d \max(\theta(P)^{-\phi(d)}, 1)$. ∎

When $d = 2$, Theorem 5 shows that QMG is optimal because we already know that for any triangulation $\mathcal{T}$, $\rho_{\mathcal{T}}(P) \geq c\theta(P)^{-1}$ by Theorem 1, and $\phi(2) = 1$.

When $d = 3$, Theorem 5 shows that QMG has an aspect ratio bound of $c\theta(P)^{-2}$, whereas the lower bound from Theorem 1 is $c\theta(P)^{-1}$. In fact, a more complicated analysis of QMG for the $d = 3$ case establishes an upper bound of $c\theta(P)^{-1}$ on $\rho_{\mathrm{QMG}}(P)$. Here is a sketch of this analysis. Recall that the only case that needs attention is the case of a large protected box B_1 next to a much smaller protected box B_2, such $\mathrm{co}(B_2) \subset \mathrm{co}(B_1)$. Suppose, for instance, that B_1 is protected in phase 0 (a similar argument applies to the case when B is protected in phase 1), and suppose that B_2 is protected in phases 1 or 2. Find the largest ancestor B^* of B_2 with the property that $\mathrm{co}(B^*) \subset \mathrm{co}(B_1)$; as above, $\mathrm{size}(B_1)/\mathrm{size}(B^*) \leq c$. Consider the chain of strongly crowded boxes from B^* down to B_2. Let P' be the intersection of P with the facet of B_1 separating it from B^*. Observe that splitting strongly crowded boxes in phases 1 and 2 of 3-dimensional QMG running applied to the chain of boxes B^* down to B_2 is very similar to phases 0 and 1 of 2-dimensional QMG running on the polygon P'. In other words, with a correct modification to the definition of $\mathrm{ex}(B)$ in two dimensions, whenever a 3D box on the boundary of B_1 is split for strong crowding, the corresponding two-dimensional box would be split for strong crowding of P'.

Since 2-dimensional QMG is optimal with respect to maximizing the minimum altitude, we conclude that $\mathrm{size}(B_2)$ is bounded below by the minimum geodesic distance δ in P' between two faces that do not meet. But now it is easy to see that $\delta/\mathrm{size}(B_1)$ is bounded above by the sharpest angle θ at $\boldsymbol{v}$. Thus, $\mathrm{size}(B_1)/\mathrm{size}(B_2) \leq c\theta(P)^{-1}$.

It is likely that this line of reasoning extends to higher dimensions, although we do not know the exact improvement to Theorem 5 possible with this analysis. Furthermore, we do not know whether our lower bound Theorem 1 on the best attainable aspect ratio is tight in dimensions higher than 3.

We can summarize the conclusions of this section with the following theorem.

Theorem 6 *Let $\mathcal{T}$ be an arbitrary triangulation of P with worst case aspect ratio denoted $\rho_{\mathcal{T}}(P)$. Then*

$$\rho_{\mathrm{QMG}}(P) \leq c_d \cdot (\rho_{\mathcal{T}}(P))^{\psi(d)}$$

where $\psi(d) = 1$ for $d = 2, 3$ and $\psi(d) \le \phi(d)$ for higher dimensions.

16 Bounded aspect ratio triangulations

In the previous section we showed that the QMG triangulation has an aspect ratio bound. In the next section we will show that, among all triangulations with bounded aspect ratio, QMG has the minimum cardinality, up to a constant factor. This requires some preliminary results that apply generally to any triangulation with bounded aspect ratio.

We have the following preliminary lemma.

Lemma 10 *Let $\mathcal{T}$ be a triangulation of a polyhedral region P whose aspect ratio is at most ρ. Let T_1, T_2 be two simplices that share a common vertex $\boldsymbol{v}$. Then $\mathrm{minalt}(T_1) \le \zeta_1(\rho, d)\,\mathrm{minalt}(T_2)$, where the function $\zeta_1(\cdot,\cdot)$ is defined below.*

We omit the proof of this lemma, which is contained in [10]. Here is a sketch. Two simplices S_1, S_2 of $\mathcal{T}$ that share an edge $\boldsymbol{v}_1\boldsymbol{v}_2$ satisfy $\mathrm{minalt}(S_1) \le \rho\,\mathrm{minalt}(S_2)$ by the chain of inequalities:

$$\begin{aligned} \mathrm{minalt}(S_1) &\le \|\boldsymbol{v}_1 - \boldsymbol{v}_2\| \\ &\le \mathrm{maxside}(S_2) \\ &= \mathrm{asp}(S_2)\,\mathrm{minalt}(S_2) \\ &\le \rho\,\mathrm{minalt}(S_2). \end{aligned} \tag{9}$$

Two simplices T_1, T_2 that share a vertex $\boldsymbol{v}$ are connected by a chain of simplices $S_1(= T_1), S_2, \ldots, S_p(= T_2)$ that all share $\boldsymbol{v}$ and such that S_i and S_{i+1} have a common edge. This is because $\mathcal{T}$ is a triangulation of P, which is a manifold with boundary. Then it can be shown that the number of simplices p that can share a common vertex is bounded above in terms of ρ because the solid angle of each S_i at $\boldsymbol{v}$ is bounded below in terms of ρ. Thus, p is bounded above in terms of ρ: it turns out that $p \le c_d\rho^{d-1}$. Thus, the lemma is true with $\zeta_1(\rho, d) = \rho^{c_d\rho^{d-1}}$.

Now for the first result of the section. This lemma bounds the rate at which simplices can grow in a bounded aspect ratio triangulation.

Lemma 11 *Let $\mathcal{T}$ be a triangulation of P whose aspect ratio bound is ρ. Let Π be a PL path in P from $\boldsymbol{x}$ to $\boldsymbol{y}$. Suppose that $\boldsymbol{x}$ is contained in a simplex*

T. *Then every simplex* T' *containing* $\boldsymbol{y}$ *satisfies*

$$\operatorname{minalt}(T') \leq c_d \zeta_1(\rho, d) \max(\operatorname{minalt}(T), \operatorname{lth}(\Pi)). \tag{10}$$

Proof. Let $F_1, \ldots, F_s$ be an enumeration of the faces of $\mathcal{T}$ met by Π. Let F_s be the lowest dimensional face of $\mathcal{T}$ containing $\boldsymbol{y}$.

Case 1, F_s *has a vertex in common with* T. Since every simplex containing $\boldsymbol{y}$ also contains F_s, then T' has a common vertex with T. In this case the lemma is true with the first term of the max in (10) by the preceding lemma.

Case 2, F_s *does not have a vertex in common with* T. In this case, define a PL continuous function $f : P \rightarrow \mathbb{R}$ that is 1 on vertices of F_s, zero on all other vertices, and is linearly interpolated by $\mathcal{T}$. Then, as in the proof of Theorem 2, there must be a simplex S that meets Π such that the gradient of f on S is at least $1/\operatorname{lth}(\Pi)$, and hence this simplex satisfies $\operatorname{minalt}(S) \leq c_d \operatorname{lth}(\Pi)$. Notice that S and F_s must have a common vertex because if not, then the gradient of f on S would be 0. Since S and T' have a common vertex, we apply the preceding lemma to bound $\operatorname{minalt}(T')$ by the second term of the max in (10), proving the lemma. ∎

Here is our other main result about bounded aspect ratio triangulations.

Lemma 12 *Let* C *be a cube in* $\mathbb{R}^d$ *of side length* s. *Let* $T_1, \ldots, T_n$ *be a set of* n *simplices with pairwise disjoint interiors, satisfying* $\operatorname{minalt}(T_i) \geq \mu$ *and* $\operatorname{asp}(T_i) \leq \rho$ *for each* i. *Suppose each* T_i *meets* C. *Then*

$$n \leq c_d \rho^d + c_d s^d / \mu^d. \tag{11}$$

Proof. For each T_i, identify a point $\boldsymbol{x}_i \in T_i \cap C$. Now contract each T_i about $\boldsymbol{x}_i$ until we obtain a new simplex T_i' such that $\operatorname{minalt}(T_i') = \mu$. Since $T_i' \subset T_i$, the set $T_1', \ldots, T_n'$ still enjoys the property that interiors are pairwise disjoint. Since $\boldsymbol{x}_i \in T_i'$, each T_i' still meets C. Finally, contraction affects maxside and minalt by the same scale factor, so $\operatorname{asp}(T_i') = \operatorname{asp}(T_i)$.

We know that $\operatorname{maxside}(T_i') = \operatorname{asp}(T_i') \operatorname{minalt}(T_i') \leq \rho\mu$. Therefore, for every point $\boldsymbol{y}$ in T_i', $\operatorname{dist}(\boldsymbol{y}, \boldsymbol{x}_i) \leq \rho\mu$. Let $\boldsymbol{x}_0$ be the centroid of C. Then we have for each i that $\operatorname{dist}(\boldsymbol{x}_i, \boldsymbol{x}_0) \leq c_d s$. Combining these inequalities yields the bound that for every point $\boldsymbol{y} \in T_i'$ for each i, $\operatorname{dist}(\boldsymbol{y}, \boldsymbol{x}_0) \leq \rho\mu + c_d s$. Thus, all of $T_1', \ldots, T_n'$ are contained in a ball B of radius $\rho\mu + c_d s$ around $\boldsymbol{x}_0$. The volume of this ball is at most $c_d(\rho^d \mu^d + s^d)$. Each simplex has volume at least $c_d \mu^d$ by (1). Since the simplices have disjoint interiors, their number is bounded above by $\operatorname{vol}(B)/(c_d \mu^d)$, which proves the lemma. ∎

17 A bound on the cardinality of QMG

In this section we show that the cardinality of the triangulation produced by QMG is always within a constant factor of optimal among all bounded aspect ratio triangulations, where the constant depends on the aspect ratio. We start with the following lemma.

Lemma 13 *Let $\boldsymbol{x}$ be an arbitrary point in P, and let T be the simplex generated by QMG that contains $\boldsymbol{x}$. (If there is more than one, the result holds for any choice of T.) Let $\mathcal{S}$ be some other triangulation of P with aspect ratio bound $\rho_{\mathcal{S}}$, and let S be the simplex in $\mathcal{S}$ that contains $\boldsymbol{x}$. (If there is more than one, then the result holds for any choice of S.) Then*

$$\mathrm{minalt}(S) \leq c_d \zeta_1(\rho_{\mathcal{S}}, d) \cdot \max(\theta^{-(2d+2)\phi(d)}, 1) \cdot \mathrm{minalt}(T). \tag{12}$$

Proof. Recall that each simplex generated by QMG is associated with a full-dimensional protected box, namely the last box in its chain. Furthermore, each full-dimensional protected box is associated with a full-dimensional anchor box as in Theorem 4. Therefore, transitively, each simplex generated by QMG is associated with an anchor box.

Let T be such a simplex, and B_a its anchor box. Clearly

$$\mathrm{minalt}(T) \leq c_d \,\mathrm{size}(B_a) \tag{13}$$

since T lies in $\mathrm{ex}(B_a)$. On the other hand, there is also an inequality in the other direction. The reason is as follow. Let B be the full-dimensional protected box containing T. Then, as argued in Section 10, $\mathrm{minalt}(T)$ is bounded below by the size of the smallest neighbor of B, which, by the proof of Theorem 5, is bounded below by $c_d \min(1, \theta(P)^{\phi(d)})\,\mathrm{size}(B)$. There is a lower bound on $\mathrm{size}(B)$ in terms of $\mathrm{size}(B_a)$ given by (4). Combining these bounds yields

$$\mathrm{minalt}(T) \geq c_d \min(1, \theta(P)^{2\phi(d)})\,\mathrm{size}(B_a). \tag{14}$$

Let $\boldsymbol{x}$ be the arbitrary point in P specified by the lemma. Let T be the simplex generated by QMG that contains $\boldsymbol{x}$, let B be the protected box associated with T, and let B_a be the anchor of B. In the next few paragraphs we will construct a PL path Σ from $\boldsymbol{x}$ to a simplex $S \in \mathcal{S}$ satisfying $\mathrm{minalt}(S) \leq \mathrm{size}(B_a)$, and such that $\mathrm{lth}(\Sigma)$ is bounded above in terms of $\mathrm{size}(B_a)$.

Observe that $\mathrm{co}(B_a)$ contains an incontractible path Π_1 by definition of "anchor." Therefore, let S_1 be the simplex in $\mathcal{S}$ that meets Π_1 and satisfies $\mathrm{minalt}(S_1) \leq c_d \,\mathrm{size}(B_a)$ as specified in Theorem 2. Let $\boldsymbol{y}$ be a point in $\Pi_1 \cap S_1$. Construct the shortest PL path (the geodesic path) from $\boldsymbol{x}$ to $\boldsymbol{y}$ lying in $\mathrm{co}(B_a)$, and call it Σ_1. Note that there is no *a priori* upper bound on $\mathrm{lth}(\Sigma_1)$ in terms of $\mathrm{size}(B_a)$ because $\mathrm{co}(B_a)$ could contain geodesic paths possibly much longer than $\mathrm{size}(B_a)$.

Let $T_1, T_2, \ldots, T_p$ be an enumeration of the QMG simplices met by Σ_1, listed in the order they are encountered starting with $\boldsymbol{x}$. Note that no simplex can appear twice in this enumeration; this is because Σ_1 is a geodesic path and therefore would not return to the same simplex more than once.

Let T_q be the first simplex in the sequence that fails to satisfy (14). If there is no such q, then take $q = p+1$. Thus, $T_1, \ldots, T_{q-1}$ all satisfy (14). We claim that $q - 1 \leq c_d \max(\theta^{-2d\phi(d)}, 1)$. This follows from (11). Observe that $T_1, \ldots, T_{q-1}$ is a set of simplices with disjoint interiors all meeting $\mathrm{ex}(B_a)$. We use (14) as a lower bound on the minimum altitudes of $T_1, \ldots, T_{q-1}$, $c_d \,\mathrm{size}(B_a)$ as the size of $\mathrm{ex}(B_a)$, and Theorem 5 to get upper bounds on aspect ratios. In this use of (11), the second term dominates the first on the right-hand side.

Suppose that $q = p + 1$, i.e., every simplex in the enumeration satisfies (14). Then define $\Sigma = \Sigma_1$; we claim that

$$\mathrm{lth}(\Sigma) \leq c_d \max(\theta^{-2d\phi(d)}, 1)\,\mathrm{size}(B_a). \tag{15}$$

This is because Σ passes through $q-1$ simplices, and the length of its segment in each simplex is at most $c_d \,\mathrm{size}(B_a)$. This choice of Σ has all the properties named above: it connects $\boldsymbol{x}$ to a point on a simplex S_1 (which satisfies $\mathrm{minalt}(S_1) \leq c_d \,\mathrm{size}(B_a)$) and satisfies (15).

The other case is that $q < p-1$. In this case, we truncate Σ_1 at the point where it enters T_q which we denote $\boldsymbol{x}_1$; call the truncated path Σ_1'. Clearly Σ_1' satisfies (15) by the same argument as in the last paragraph. Now notice that the anchor box for T_q cannot be B_a because T_q does not satisfy (14) by choice of q.

Therefore, identify the anchor of T_q, which we will call B_b, and start the construction anew from $\boldsymbol{x}_1$. In other words, find the incontractible path Π_2 in $\mathrm{co}(B_b)$, find the simplex S_2 of $\mathcal{S}$ that meets the incontractible path and has altitude at most $c_d \,\mathrm{size}(B_b)$ and let Σ_2 be the path $\boldsymbol{x}_1$ to a point in $S_2 \cap \Pi_2$. Note that $\mathrm{size}(B_b) \leq \mathrm{size}(B_a)/2$, because B_b must be smaller than B_a so that (14) can be satisfied for the new anchor. Find the first simplex in this new path that fails to satisfy (14) for B_b and so on. Notice that Σ_2', the truncation of Σ_2, satisfies (15) with $\mathrm{size}(B_b)$ taking the place of $\mathrm{size}(B_a)$ on the right-hand side. Therefore, the upper bounds given by the right-hand side of (15) on $\mathrm{lth}(\Sigma_1'), \mathrm{lth}(\Sigma_2'), \ldots$ form a series decreasing by a factor of 2 each time. Eventually the procedure terminates at Σ_l because there is a finite lower bound on the smallest protected box in QMG. When the procedure terminates, concatenate $\Sigma_1', \Sigma_2', \ldots, \Sigma_{l-1}', \Sigma_l$ into a PL path Σ. (This concatenation is possible because Σ_1' ends at $\boldsymbol{x}_1$, which is where Σ_2' begins, and so on.)

This path Σ has the following properties. It satisfies (15) with the original B_a on the right-hand side, multiplied by an additional factor of 2 that arises from summing a decreasing geometric series. It connects $\boldsymbol{x}$, the given point in P contained in a simplex T of QMG anchored at B_a, to a point $\boldsymbol{y}$ that is in a simplex S_l in triangulation $\mathcal{S}$ and that satisfies $\mathrm{minalt}(S_l) \leq c_d \,\mathrm{size}(B_a)$.

This is exactly the setup we need to apply Lemma 11. Let S be the simplex in $\mathcal{S}$ that contains $\boldsymbol{x}$. From Lemma 11 applied to Σ we conclude that

$$\begin{aligned}\mathrm{minalt}(S) &\leq c_d\zeta_1(\rho_{\mathcal{S}}, d) \cdot \max(\mathrm{lth}(\Sigma), \mathrm{minalt}(S_l)) \\ &\leq c_d\zeta_1(\rho_{\mathcal{S}}, d) \cdot \max(\theta^{-2d\phi(d)}, 1)\,\mathrm{size}(B_a). \qquad (16)\end{aligned}$$

The second line was obtained by substituting the bound (15) for $\mathrm{lth}(\Sigma)$, and then noting that this bound dominates the upper bound of $c_d \,\mathrm{size}(B_a)$ that applies to $\mathrm{minalt}(S_l)$.

Now finally, the lemma is proved, because we combine (16) with the bound on $\mathrm{size}(B_a)$ in terms of $\mathrm{minalt}(T)$ given by (14). ∎

Theorem 7 *Let $n_{\mathrm{QMG}}(P)$ be the number of simplices in the triangulation produced by QMG. Let $\mathcal{S}$ be some other triangulation of P with aspect ratio bound $\rho_{\mathcal{S}}$, and let the cardinality of $\mathcal{S}$ be $n_{\mathcal{S}}$. Then*

$$n_{\mathrm{QMG}}(P) \leq c_d\zeta_1(\rho_{\mathcal{S}}, d)^d \rho_{\mathcal{S}}^d \cdot \max(\theta^{-(2d+2)d\phi(d)}, 1) \cdot n_{\mathcal{S}}. \qquad (17)$$

Proof. Let $f_{\mathrm{QMG}} : P \to \mathbb{R}$ be the piecewise constant function defined as follows. Let T be a simplex generated by QMG. The value of f_{QMG} on T is defined to be $1/\mathrm{vol}(T)$. On boundaries of simplices, a measure-zero set, we leave f_{QMG} undefined. Function $f_{\mathcal{S}} : P \rightarrow \mathbb{R}$ is defined similarly in terms of $\mathcal{S}$. Note that

$$n_{\text{QMG}} = \int_{\boldsymbol{x} \in P} f_{\text{QMG}}(\boldsymbol{x})\, d\boldsymbol{x}$$

because the value of the integral over each individual simplex is exactly 1. A similar expression holds for $n_{\mathcal{S}}$.

Define piecewise constant functions $g_{\text{QMG}} : P \rightarrow \mathbb{R}$ to be $1/\operatorname{minalt}(T)^d$ on T, where T is a simplex in QMG, $g_{\mathcal{S}}$ similarly for $\mathcal{S}$. Finally, define $h_{\mathcal{S}}$ to the piecewise constant function that is $1/\operatorname{maxside}(S)^d$ on S, as S ranges over simplices in $\mathcal{S}$.

Then we have the following chain of inequalities.

$$\begin{aligned}
n_{\text{QMG}} &= \int_{\boldsymbol{x} \in P} f_{\text{QMG}}(\boldsymbol{x})\, d\boldsymbol{x} \\
&\leq c_d \int_{\boldsymbol{x} \in P} g_{\text{QMG}}(\boldsymbol{x})\, d\boldsymbol{x} \\
&\leq c_d \zeta_1(\rho_{\mathcal{S}}, d)^d \cdot \max(\theta^{-(2d+2)d\phi(d)}, 1) \cdot \int_{\boldsymbol{x} \in P} g_{\mathcal{S}}(\boldsymbol{x})\, d\boldsymbol{x} \\
&\leq c_d \zeta_1(\rho_{\mathcal{S}}, d)^d \rho_{\mathcal{S}}^d \cdot \max(\theta^{-(2d+2)d\phi(d)}, 1) \cdot \int_{\boldsymbol{x} \in P} h_{\mathcal{S}}(\boldsymbol{x})\, d\boldsymbol{x} \\
&\leq c_d \zeta_1(\rho_{\mathcal{S}}, d)^d \rho_{\mathcal{S}}^d \cdot \max(\theta^{-(2d+2)d\phi(d)}, 1) \cdot \int_{\boldsymbol{x} \in P} f_{\mathcal{S}}(\boldsymbol{x})\, d\boldsymbol{x} \\
&= c_d \zeta_1(\rho_{\mathcal{S}}, d)^d \rho_{\mathcal{S}}^d \cdot \max(\theta^{-(2d+2)d\phi(d)}, 1) \cdot n_{\mathcal{S}}.
\end{aligned}$$

In these inequalities, we used (1) to obtain the second line, (12) for the third line, the aspect ratio bound for $\mathcal{S}$ for the fourth line, and (1) again for the fifth line. This proves the theorem. ∎

Note that this theorem allows the ratio $n_{\text{QMG}}/n_{\mathcal{S}}$ to be arbitrarily large if the competing triangulation $\mathcal{S}$ has bad aspect ratio. This is not merely an artifact of our analysis but is actually a feature of bounded aspect-ratio triangulations, as illustrated by the following example. Consider a $p \times 1$ rectangle with $p \gg 1$. On such a domain, QMG would require $O(p)$ triangles (as would any algorithm guaranteeing bounded-aspect ratio), but this domain can be triangulated with just two triangles by inserting a diagonal. This latter triangulation has aspect ratio of $\Omega(p)$.

Note also that $\rho_{\mathcal{S}} \geq c_d \theta^{-1}$ by Theorem 1. Thus, the entire right-hand side of (17) can be bounded above with the more compact formula $f(\rho_{\mathcal{S}}, d) \cdot n_{\mathcal{S}}$.

18 Running time analysis

In this section we briefly discuss the running time of QMG. The running time for the separation stages is proportional to the number of boxes created multiplied by the time per box. There is no prior upper bound on the number of boxes in terms of the input. There is also no prior upper bound on the number of boxes in terms of the output, that is, in terms of the number of simplices produced, which we denote s. However, a modification to QMG would allow us to claim that the total number of boxes is bounded above by a multiple of s. The modification would be an additional operation to short-circuit a series of splitting operations that make no progress. More specifically, the modification is as follows. When we split a box, we check if only one of its children has nonempty content. If so, we discard the other children, and we immediately shrink that box by a power-of-two factor, until it is sufficiently small that we can be guaranteed that the next split will produce more than one child with nonempty content.

The amount of time to process a box depends on the combinatorial complexity of its content. A crude upper bound is that the complexity of the content is bounded by $O(N)$, where N is the complexity of the input domain P. Processing the content requires a connected component computation; the time for this computation in higher dimensions is $O(N^2)$, although, as mentioned in Section 5, more efficient algorithms are available for 2 and 3 dimensions.

Thus, an estimate for the separation stage running time, using the modification mentioned above, is $O(N^2 s)$ operations. The operations in the alignment stage (checking complete coverage) can be done with a hash table as mentioned in Section 7. Thus, alignment requires $O(s)$ operations. Finally, the triangulation part of the algorithm also requires $O(s)$ operations. Thus, the total running time is $O(N^2 s)$.

19 Implementation

A two-dimensional version of QMG called "tripoint" was implemented by S. Mitchell in C++ and is available on the web [11]. A full version of QMG has been implemented in C++ by S. Vavasis; the implementation QMG1.1 is more general than the version described in this paper because it can also handle nonmanifold features, including several kinds of internal boundaries.

QMG1.1 is available on the Worldwide Web [16] and has a number of users in several countries. It is slated to be the main mesh generator in a future release of //Ellpack. The implementation is slightly different from the algorithm described in this paper; in particular, the alignment procedure uses an adaptive method for selecting tolerances, and a different rule for choosing a close face. Computational experiments will be described elsewhere.

20 Open questions

Some of the open questions raised by this work include:

1. Is there a triangulation algorithm with stronger optimality properties? For instance, the QMG aspect ratio is optimal (up to a constant factor) only in dimensions 2 and 3.

2. Several open questions were posed in Section 14. For instance, is there a characterization of the maximal value of the min-altitude in a triangulation of 3D polyhedra?

3. Can this work be extended to curved boundaries? It appears that the main bottleneck is a solution to the subproblem of triangulating a uniform grid of boxes posed in the companion paper [12].

4. Is there a mesh generation algorithm for three-dimensional domains that guarantees dihedral angles bounded by $\pi/2$? Such a bound is important for some finite element problems [15]. It is known [3] how to solve the corresponding problem in two dimensions.

5. Can the running time bound be improved?

6. The optimality properties demonstrated here all involve constant factors which are apparently very large. The QMG implementation in practice often produces meshes that are factors of 50 off from minimum number of tetrahedra, and also factors of 50 off from the best aspect ratio. This leads to the following practical open question: Is there a mesh generation algorithm with theoretical guarantees comparable to QMG's, and yet with better theoretical constants or better performance in practice?

21 Acknowledgments

Preparation of this paper has been an on-again, off-again process spanning a number of years. Several people have given us very useful comments on earlier versions including Chandrajit Bajaj, Marshall Bern, Paul Chew, John Gilbert, Joe Mitchell, and the two referees of [13].

References

[1] B. S. Baker, E. Gross, and C. Rafferty. Nonobtuse triangulation of polygons. *Discrete and Comp. Geometry*, 3:147–168, 1988.

[2] M. Bern, D. Dobkin, and D. Eppstein. Triangulating polygons without large angles. Preprint, 1991.

[3] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. In D. Z. Du and F. K. Hwang, editors, *Computing in Euclidean Geometry*, pages 47–123. World Scientific, Singapore, 1995.

[4] M. Bern, D. Eppstein, and J. Gilbert. Provably good mesh generation. *J. Comp. Sytem Sciences*, pages 384–409, 1994.

[5] M. Bern and P. Plassmann. Mesh generation. Preprint, 1996.

[6] B. Chazelle and L. Palios. Triangulating a nonconvex polytope. In *Proc. 5th ACM Symposium on Computational Geometry*, pages 393–400. ACM Press, 1989.

[7] L. P. Chew. Guaranteed-quality triangular meshes. Technical Report TR-89-983, Department of Computer Science, Cornell University, 1989.

[8] L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. In *Proc. 9th ACM Symposium on Computational Geometry*, pages 274–280. ACM Press, 1993.

[9] Claes Johnson. *Numerical Solution of Partial Differential Equations by the Finite Element Method*. Cambridge University Press, Cambridge, 1987.

[10] G. Miller, W. Thurston, S.-H. Teng, and S. Vavasis. Geometric separators for finite element meshes. To appear in *SIAM J. Scientific Computing*, 1996.

[11] S. Mitchell. Tripoint information. Available at `http: // sass577.endo.sandia.gov:80 / 9225 / Personnel / samitch / csstuff / csguide.html` , 1996.

[12] S. A. Mitchell and S. Vavasis. An aspect ratio bound for triangulating a d-grid cut by a hyperplane (extended abstract). In *Proc. 12th ACM Symposium on Computational Geometry*, pages 48–57, 1996. A full version is available on line at `ftp: // ftp.cs.cornell.edu / pub / vavasis / papers / asp9.ps.Z`.

[13] S. A. Mitchell and S. A. Vavasis. Quality mesh generation in three dimensions. In *Proceedings of the ACM Computational Geometry Conference*, pages 212–221, 1992. Also appeared as Cornell C.S. TR 92-1267.

[14] J. Ruppert. A new and simple algorithm for quality 2-dimensional mesh generation. In *Proc. 4th ACM-SIAM Symp. on Disc. Algorithms*, pages 83–92. ACM Press, 1993.

[15] S. A. Vavasis. Stable finite elements for problems with wild coefficients. Technical Report 93–1364, Department of Computer Science, Cornell University, Ithaca, NY, 1993. To appear in *SIAM J. Numerical Analysis*.

[16] S. A. Vavasis. QMG: Software for finite-element mesh generation. See `http: // www.cs.cornell.edu / home / vavasis / qmg-home.html`, 1995.