An aspect ratio bound for triangulating a d-grid cut by a hyperplane

Scott A. Mitchell* Stephen A. Vavasis†

August 14, 1995

Abstract

We consider the problem of triangulating a d-dimensional uniform grid of d-cubes that is cut by a k-dimensional affine subspace. The goal is to obtain a triangulation with bounded aspect ratio. To achieve this goal, we allow some of the box faces near the affine subspace to be displaced. This problem has applications to finite element mesh generation. For general d and k, the bound on aspect ratio that we attain is double-exponential in d. For the important special case of $d = 3$, the aspect ratio bound is small enough that the technique is useful in practice.

1 Introduction

Recently the authors have developed the QMG mesh generator [10, 12], which is an algorithm for finite-element mesh generation of polyhedral regions in d dimensions. The main feature of the QMG mesh generator is that it guarantees an upper bound on the aspect ratio of all simplices occurring in the final mesh. The **aspect ratio** of a simplex is defined to be its maximum side-length divided by its minimum altitude. Small aspect ratio is important for accuracy in the finite element method [7].

The QMG algorithm generates boxes covering boundary faces of the polyhedral region P; these boxes are grouped into small clusters. Each such cluster contains boxes all of the same size and all the boxes in a cluster are associated with exactly one k-dimensional affine subspace, which is a face of P. Thus, QMG gives rise to the subproblem of triangulating equally sized boxes cut by a hyperplane.

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

†Department of Computer Science, Cornell University, Ithaca, NY 14853, vavasis@cs.cornell.edu.
This author's work was supported in part by an NSF Presidential Young Investigator grant with matching funds from Xerox, AT&T, Sun and Tektronix. This work was partially supported by Xerox Corporation, during a visit to the Xerox Palo Alto Research Center. This work was partially supported by the Research Institute for Advanced Computer Science (RIACS) during a visit to RIACS, sponsored by NASA under contract NAS 2-13721.

In this paper we consider this subproblem, namely, triangulating a uniform mesh of boxes in $\mathbf{R}^d$ that is cut by a k-dimensional affine subspace denoted F, where k lies between 0 and $d-1$. We would like the triangulation to have the following properties:

1. The triangulation is a valid simplicial complex, i.e., the intersection of any pair of simplices is a simplicial subface of both, and the union of the simplices is $\mathbf{R}^d$.

2. The triangulation respects the affine subspace. In other words, the intersection of F with any simplex T in the triangulation must be a simplicial subface of T.

3. The triangulation respects the mesh. That is, the intersection of any box B with T must be a simplicial subface of T.

4. The triangulation has bounded aspect ratio. This means that there is an upper bound in terms of d and k only on the aspect ratio of all simplices.

5. Every vertex of the triangulation is either a vertex of a box or lies on F.

In general this problem may not have a solution because F may pass very close to a subface of a box. Therefore, we relax requirement 3 by allowing "warping" of boxes. We allow box faces close to F to be displaced. The amount of displacement must be bounded by a constant fraction of the box's side length.

For example, consider Figure 1 in which $d=2$ and $k=1$. Observe that the line F passes very close to the box vertices at $(0,1)$ and $(1,2)$, and somewhat farther from $(2,3)$ and $(3,4)$, etc. Accordingly, we allow $(0,1)$, $(1,2)$ and $(2,3)$ to be displaced, and the boxes that contain them to be deformed in the course of triangulation.

For the case $d=2$ with $k=0,1$, warping rules were proposed by Bern, Eppstein and Gilbert [3] that solve this problem. In our earlier work [9] on mesh generation, we partially analyzed a warping rule for $d=3$ using an argument based on cases. A case-based construction without an aspect ratio analysis was provided for $d=3$ by Frey et al [6]. In this paper, we provide a complete analysis of the problem for all d and k.

A different approach to guaranteed aspect ratio, which appears to be hard to generalize to dimensions higher than $d=2$, was introduced by Chew [5].

2 Algorithm description

To simplify notation, let us assume that the grid of boxes is the integer lattice; in other words, each box is of the form $[s_1,s_1+1]\times\cdots\times[s_d,s_d+1]$ where $s_1,\ldots,s_d$ are all integers. A *box face* is a face of one of these boxes. A box face B can always be written as $[s_1,t_1]\times\cdots\times[s_d,t_d]$ where for each i, s_i is an integer and t_i is equal to either s_i or s_i+1. We say that B *extends* over dimension i if the ith factor in this cross product is $[s_i,s_i+1]$ (rather than $[s_i,s_i]$). The *dimension* of B, denoted $\dim(B)$, is the number of dimensions over which B extends.

Our algorithm has two parts: preprocessing and triangulation. In the preprocessing phase, we identify the "close point" for each box face B. This close point of $\mathbf{R}^d$ is a point lying on F that is close to B.

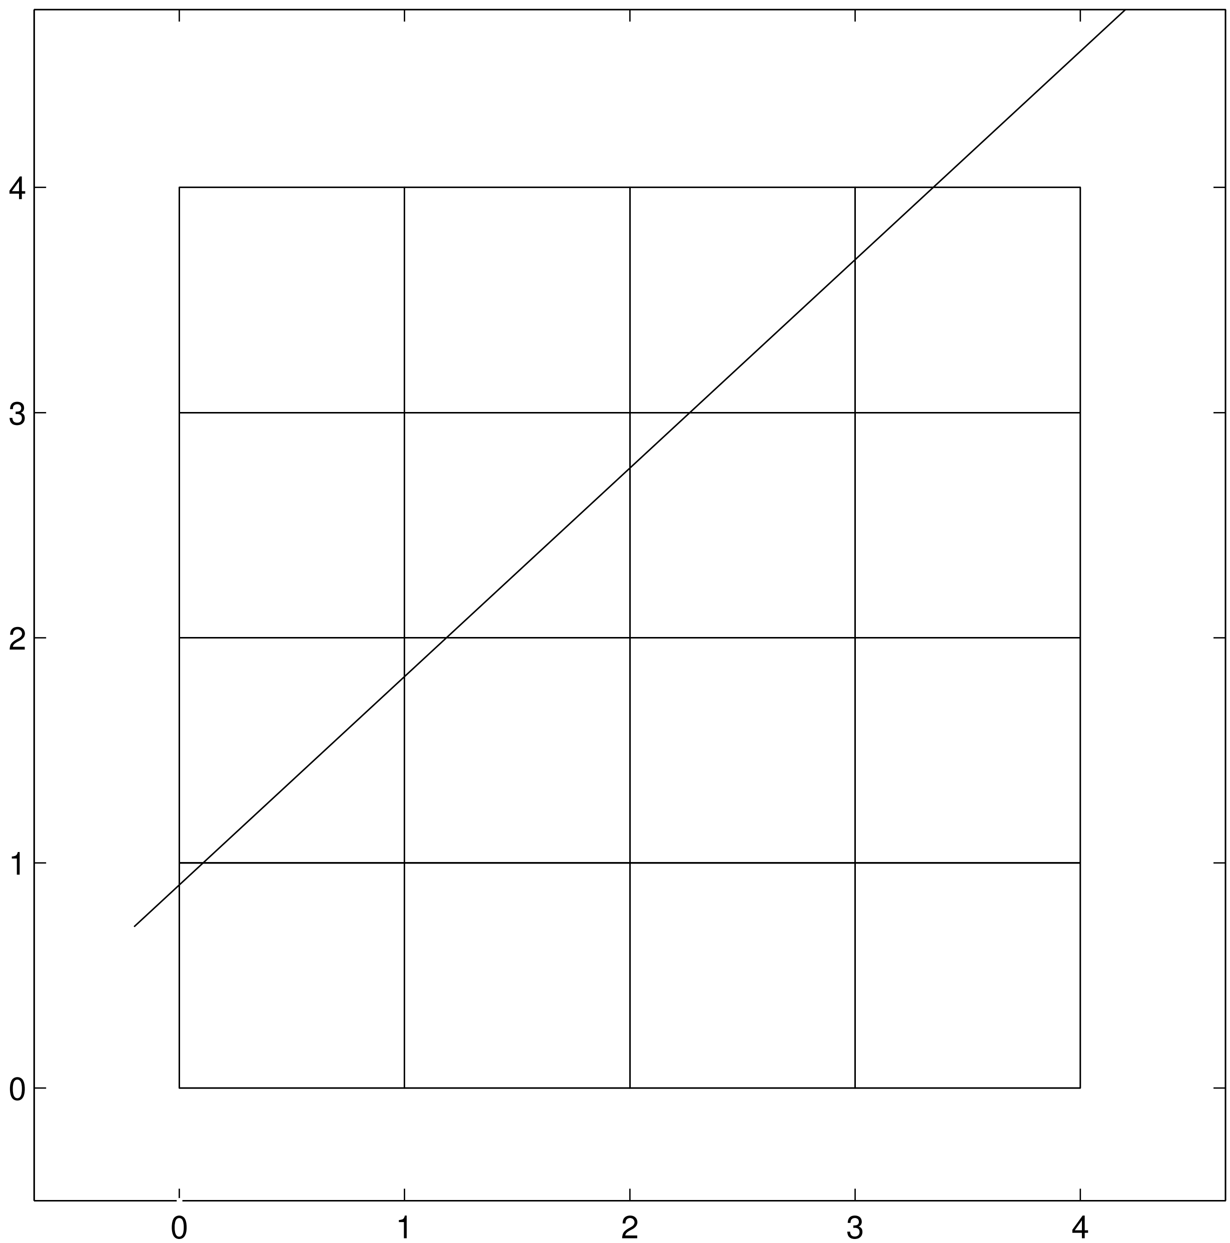

Figure 1: In the $d = 2, k = 1$ case, the goal is to triangulate the mesh and the line crossing it while maintaining bounded aspect ratio. The result of our algorithm applied to this input is illustrated in Figure 2.

We let $\mathrm{cl}(B)$ denote the close point of B. Not all box faces have a close point; for a box face with no close point, we write $\mathrm{cl}(B) = \emptyset$. The rules for selecting a close point are as follows.

First, assign a global priority order (a total order) to all box faces in the grid. Dimension-0 box faces (vertices) must have highest priority, dimension-1 box faces next priority, and so on. Within a given dimension, the priority order can be arbitrary; for instance, lexicographic order on the coordinate entries is suitable.

Next, we introduce parameters $\epsilon_0, \ldots, \epsilon_d$ that specify how near F must pass to B in order to be considered "close" to B. We use a distinct parameter for each dimension, i.e., parameter ϵ_r is used whenever $\dim(B) = r$. These parameters must satisfy several inequalities including the following:

$$0 \leq \epsilon_d \leq \cdots \leq \epsilon_1 \leq \epsilon_0 < 0.5.$$

Further requirements on these parameters will be imposed below. In Section 7 we give the precise formula in terms of d and k for defining these parameters.

Let B be a box face of dimension i, where $0 \leq i \leq d$. Define $N(B)$ to be the infinity-norm neighborhood of size ϵ_i around B; in other words,

$$N(B) = \{x : \text{there exists } y \in B \text{ such that } \|x - y\|_\infty \leq \epsilon_i\}.$$

Thus, for any B, $N(B)$ is an axis-parallel parallelepiped. The rules for choosing the close point of B to F are as follows. We first assign close points to all dimension-0 faces, then dimension 1, and so on up through dimension d. The close point for a dimension-i face is determined by close points of its subfaces of dimension less than i if any exist.

(R1) If B has any proper subface with a close point, then among all such subfaces, choose the one B' with highest priority and let $\mathrm{cl}(B) := \mathrm{cl}(B')$.

(R2) If (R1) fails, but the intersection of F with $N(B)$ is nonempty, let $\mathrm{cl}(B)$ be the point in $F \cap N(B)$ closest to B. Here "closest" refers to infinity-norm distance. Break ties arbitrarily.

(R3) If (R1) and (R2) fail, then $\mathrm{cl}(B) := \emptyset$.

Notice that closeness is inductive in the following sense. If B, B' are two box faces with $B \subset B'$ and $\mathrm{cl}(B) \neq \emptyset$, then $\mathrm{cl}(B') \neq \emptyset$. Notice also that if the close point of B is determined by rule (R1), and the particular subface with highest priority in rule (R1) is denoted C, then the close subface of C must have been determined by rule (R2). The reason is that if $\mathrm{cl}(C)$ had been determined by (R1) then C would have had a proper subface with a close point, and this proper subface would be a higher priority subface of B than C since lower dimensions have higher priority.

Now, with each box subface B we associate a point $\overline{\mathrm{cl}}(B)$ defined as follows. If $\mathrm{cl}(B) \neq \emptyset$, we let $\overline{\mathrm{cl}}(B) := \mathrm{cl}(B)$. Else if $\mathrm{cl}(B) = \emptyset$, we let $\overline{\mathrm{cl}}(B)$ be highest-priority dimension-0 subface of B. Thus, in all cases $\overline{\mathrm{cl}}(B)$ is a point in $\mathbf{R}^d$ on or near B.

Now, finally, we explain the triangulation algorithm. Let a *chain* be a sequence of $d+1$ nested box faces $B_0 \subset B_1 \subset \cdots \subset B_d$, where for each i, $\dim(B_i) = i$. Make a simplex that is the convex hull of $\overline{\text{cl}}(B_0), \ldots, \overline{\text{cl}}(B_d)$ and include this in the triangulation. Thus, the triangulation consists of all simplices induced by chains. The exception is the case of a chain for which $\overline{\text{cl}}(B_i) = \overline{\text{cl}}(B_j)$ for some indices $i \neq j$. In this case the simplex associated with the chain is null and may be dropped. Notice that in the case of a box with no close subfaces, the rule in the last paragraph yields "Kuhn's triangulation" [11] in which precisely $d!$ of the possible chains are non-null per box.

For example, in Figure 2 we have illustrated the case $d = 2, k = 1$ with $\epsilon_0 = 0.15$ and $\epsilon_1 = \epsilon_2 = 0$. (The motivation for these values of $\epsilon_0, \epsilon_1, \epsilon_2$ is detailed in upcoming sections.) The infinity-norm neighborhood around each vertex is illustrated with a square. Notice that the line passes through the neighborhood of vertex $(1, 2)$. Thus, this vertex has a close point, as does every box edge and square that contains it. For instance, the square $[0, 1] \times [1, 2]$ has this vertex as a subface, and in fact, it is the highest-priority subface of this box with a close point. (The priority rule used in this figure favors objects with higher-numbered coordinates.) An example of a non-null chain in this box would be vertex $(1, 1)$ (with no close point), edge $[0, 1] \times \{1\}$ (with close point inherited from $(0, 1)$), and finally the square $[0, 1] \times [1, 2]$ (with close point inherited from $(1, 2)$).

This concludes the description of our algorithm, except that we still need to select the parameters $\epsilon_0, \ldots, \epsilon_d$. In the upcoming sections we describe the rules for these parameters and analyze the aspect ratio of our construction.

3 Some of the ϵ_i's are zero

In this section we demonstrate two properties of (R1) and (R2) that, without loss of generality, allow us to choose $\epsilon_{d-k} = \epsilon_{d-k+1} = \cdots = \epsilon_d = 0$. Recall that k stands for the dimension of F and lies between 0 and $d - 1$.

Lemma 1 *Let B be a box face of dimension m such that $m > d - k$. Then rule (R2) can never apply to B.*

Remark. To show that rule (R2) never applies, we must prove that if F meets $N(B)$, then F must also meet $N(B')$ for some proper subface B' of B (so that (R1) would take precedence). Consider Figure 2; consider specifically the intersection of F with a 2-dimensional box like $[0, 1] \times [1, 2]$. Notice that F cannot meet the neighborhood of this square unless it first penetrates the neighborhood of one of the vertices or edges of that square, which means that (R1) would take precedence. This is the gist of the following proof.

Proof. Suppose F meets $N(B)$. Since B is a box face, we can write it as a product of intervals $[s_1, t_1] \times \cdots \times [s_d, t_d]$, where for each i, s_i and t_i are integers. Furthermore, for each i, either $s_i = t_i$ or $s_1 + 1 = t_i$; let the set of indices in the latter category be denoted as J, that is, J is the set of indices over which B extends. Thus, $|J| = m$.

By definition, we know that

$$N(B) = [s_1 - \epsilon_m, t_1 + \epsilon_m] \times \cdots \times [s_d - \epsilon_m, t_d + \epsilon_m].$$

Since F meets $N(B)$, F is affine, and $N(B)$ is a convex polytope, then their intersection $F \cap N(B)$ is a nonempty convex polytope embedded in F. This polytope must have an

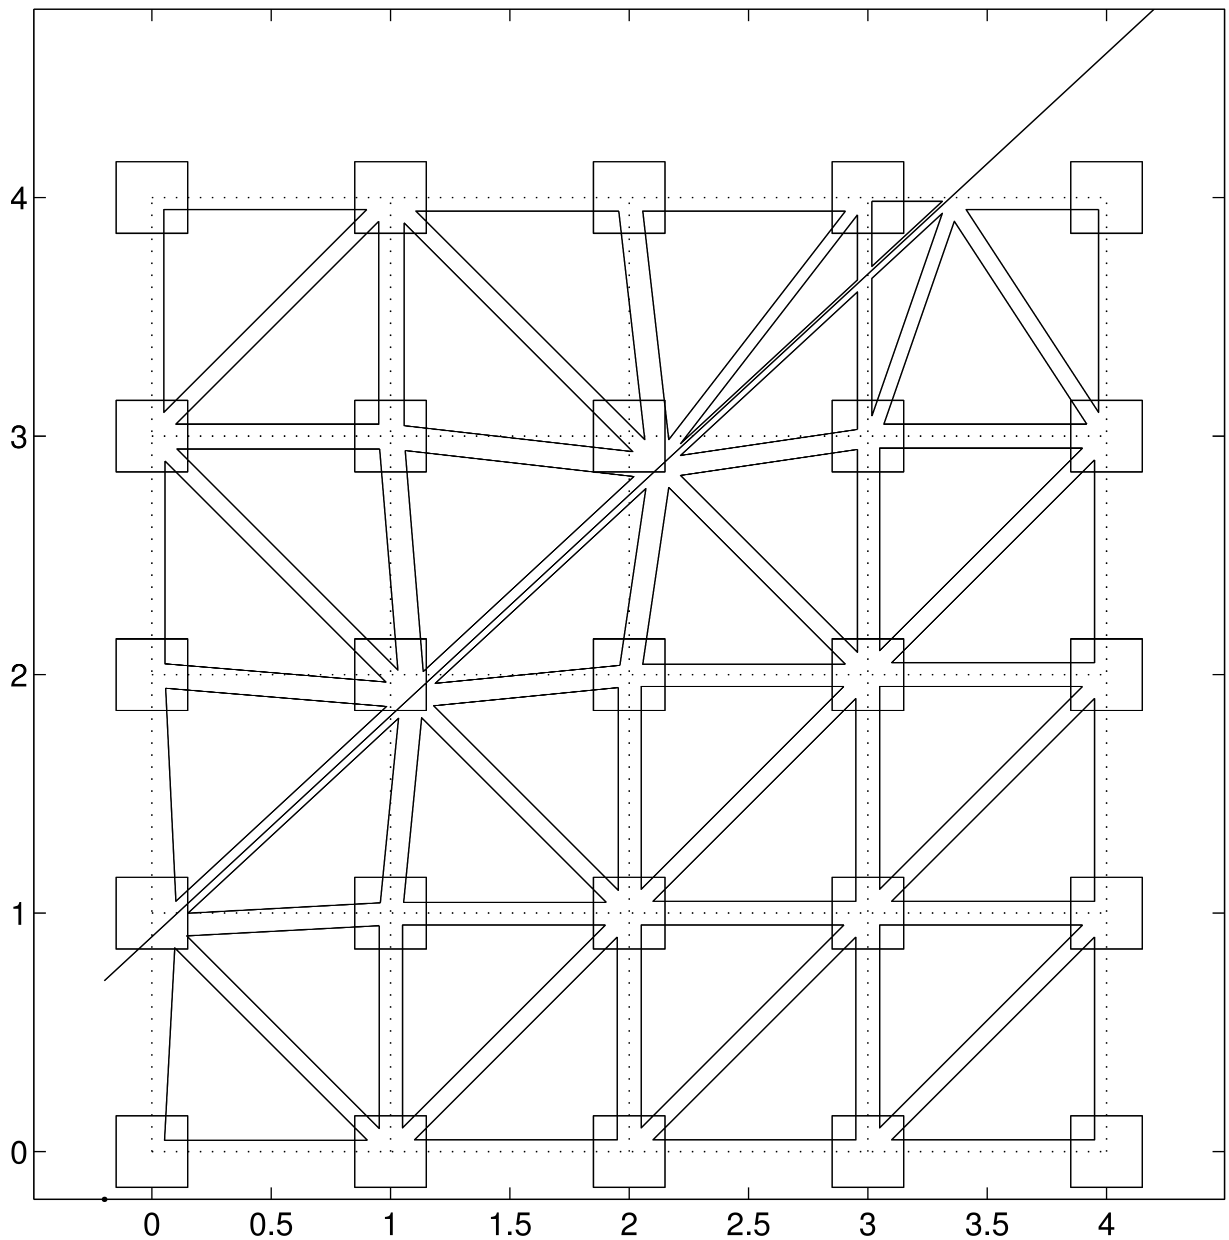

Figure 2: The result of our algorithm applied to the geometry displayed in Figure 1. The dotted lines are the original grid. The neighborhoods around vertices are drawn as small squares. Each triangle of the triangulation has been slightly contracted in preparing this figure.

extreme point, say x, and x must be the intersection of a face on the $d-k$-dimensional skeleton of $N(B)$ with F. To simplify notation, let us suppose that this intersection point x lies on the following $d-k$ face G of $N(B)$:

$$G = \{s_1 - \epsilon_m\} \times \{s_2 - \epsilon_m\} \times \cdots \times \{s_k - \epsilon_m\} \times [s_{k+1} - \epsilon_m, t_{k+1} + \epsilon_m] \times \cdots \times [s_d - \epsilon_m, t_d + \epsilon_m].$$

Since $m > d-k$, there must be at least one index $j \in \{1, \ldots, k\}$ that also lies in J. Consider the subface B' of B such that we replace $[s_j, t_j]$ in the crossproduct of B by $\{s_j\}$. This is a subface of B of dimension $m-1$ containing G. Since $\epsilon_{m-1} \geq \epsilon_m$, we see that G defined above must be a subset of $N(B')$. Therefore, F meets $N(B')$. This shows that (R1) would take precedence, so (R2) could not be applied to B. ∎

The preceding lemma shows that, without loss of generality, we may take $\epsilon_{d-k+1} = \cdots = \epsilon_d = 0$.

Lemma 2 *Let B be a box face of dimension exactly $d-k$ and suppose that rule (R2) above applies to B. Then the close point of B in rule (R2) lies in B, i.e., its distance from B is zero.*

Remark. Consider the example in Figure 2. Consider the intersection of F with a neighborhood of a box edge like $\{3\} \times [3, 4]$. If F meets the neighborhood of this edge, but misses the neighborhoods of its two subfaces, then it must meet the edge itself. This is the gist of this lemma.

Proof. To simplify notation, let us assume that B extends over the last $d-k$ dimensions so that we can write it as

$$B = \{s_1\} \times \cdots \times \{s_k\} \times [s_{k+1}, s_{k+1} + 1] \times \cdots \times [s_d, s_d + 1].$$

Let U be the affine hull of B, that is, the affine subspace of $\mathbf{R}^d$ given by $\{s_1\} \times \cdots \times \{s_k\} \times \mathbf{R}^{d-k}$. Since F has dimension k, in the nondegenerate case F meets U at exactly one point x. We will assume nondegeneracy for now, and the degenerate case is considered below. If $x \in B$ then we are done. So suppose $x \notin B$; we will derive a contradiction.

We claim that $x \notin B$ implies $x \notin N(B)$. To see this suppose $x \in N(B)$, i.e., for $i = k+1, \ldots, d$, $x_i \in [s_i - \epsilon_{d-k}, s_i + 1 + \epsilon_{d-k}]$. (And we already know $x_i = s_i$ for $i = 1, \ldots, k$ since $x \in U$.) Because $x \notin B$, there is at least one $i \in \{k+1, \ldots, d\}$ such that $x_i \notin [s_i, s_i + 1]$. For instance, suppose $x_{k+1} < s_{k+1}$ so that $x_{k+1} \in [s_{k+1} - \epsilon_{d-k}, s_{k+1}]$. Then the bounds derived on the entries of x in the three preceding sentences imply that $x \in N(B')$ where B' is the subface of B given by

$$B' = \{s_1\} \times \cdots \times \{s_{k+1}\} \times [s_{k+2}, s_{k+2} + 1] \times \cdots \times [s_d, s_d + 1].$$

This is because $\epsilon_{d-k-1} \geq \epsilon_{d-k}$. But this contradicts the use of (R2) to find the close point of B. Thus, $x \notin N(B)$.

Since (R2) applies to B, there is some point $y \in F \cap N(B)$. Consider the line segment from x to y. Since $y \in N(B)$ and $x \notin N(B)$, there must be a point z on the segment such that $z \in N(B)$ but any point strictly between z and x is not in $N(B)$. This point z lies on F because it is a convex combination of x and y. Furthermore, for each $i = 1, \ldots, k$, we know that $x_i = s_i$ and $y_i \in [s_i - \epsilon_{d-k}, s_i + \epsilon_{d-k}]$. Therefore, every point on the segment from x to y, including z remains in the the interval $[s_i - \epsilon_{d-k}, s_i + \epsilon_{d-k}]$ for all coordinate entries $i = 1, \ldots, k$. Therefore, since z is a boundary crossing for $N(B)$, there must be some coordinate $i \notin \{1, \ldots, k\}$, such that z_i is at the boundary of $[s_i - \epsilon_{d-k}, s_i + 1 + \epsilon_{d-k}]$: for instance, say that $z_{k+1} = s_{k+1} - \epsilon_{d-k}$. But one checks that $z \in N(B')$ where B' is the subface of B given above. Again, this contradicts the use of (R2).

Thus, we see that the assumption $x \notin B$ leads to a contradiction, so $x \in B$.

We now consider the degenerate case, which also leads to a contradiction. Let y be the close point in $F \cap N(B)$ determined by (R2). Because of the degenerate intersection of F and U, there exists a nonzero vector $h \in \mathbf{R}^d$ such that the first k entries of h are zeros (i.e., h is parallel to U) and such that $y + \lambda h \in F$ for all scalar values of λ (i.e., h is parallel to F). Since $N(B)$ is bounded, for sufficiently large λ, $y + \lambda h \notin N(B)$. Thus, there is a choice of $\bar{\lambda}$ such that $y + \bar{\lambda} h \in N(B)$ but $y + \lambda h \notin N(B)$ for all $\lambda > \bar{\lambda}$. Since h has zeros in positions $1, \ldots, k$, we know that $y + \bar{\lambda} h$ must be at a boundary value of $[s_i - \epsilon_{d-k}, s_i + 1 + \epsilon_{d-k}]$ for some coordinate i not in $\{1, \ldots, k\}$. But then, arguing as above, $y + \bar{\lambda} h$ would have to lie in $N(B')$ for some subface B' of B, a contradiction to the use of (R2). ∎

The preceding lemma shows that we can take $\epsilon_{d-k} = 0$ as well.

4 Aspect ratio in terms of a determinant

As mentioned in the introduction, the aspect ratio of a simplex is defined as its maximum side-length divided by its minimum altitude. Since all of the simplices in our triangulation lie inside a unit cube or a slightly distorted unit cube, the maximum side length is easily seen to be at most $(1 + 2\epsilon_0)\sqrt{d}$. (In this section, distances are measured in the 2-norm.) Therefore, the aspect ratio is bounded above provided that we can show a lower bound on the minimum altitude.

The purpose of this section is to obtain a lower bound on the minimum altitude in terms of a determinant. If S is a simplex whose vertices are $v_0 \ldots, v_d$, let

$$\Delta = \left| \det \begin{bmatrix} v_1 - v_0 \\ \vdots \\ v_d - v_0 \end{bmatrix} \right|. \tag{1}$$

Although it appears that v_0 plays a special role in (1), in fact Δ is invariant under renumbering of the v_i's. Indeed, Δ is precisely $d!$ multiplied by the volume of S.

Lemma 3 *Let S be a d-simplex embedded in $\mathbf{R}^d$ whose maximum side-length is s, and let Δ be defined by* (1). *Then the minimum altitude of S is at least Δ/s^{d-1}.*

Proof. Let w be the altitude vector from v_d to the plane spanned by $v_0, \ldots, v_{d-1}$, and without loss of generality, suppose it is the smallest.

Then, since w is orthogonal to this plane, it satisfies the $d-1$ equations $\langle w, v_1 - v_0 \rangle = 0$, $\langle w, v_2 - v_0 \rangle = 0, \ldots, \langle w, v_{d-1} - v_0 \rangle = 0$, where $\langle a, b \rangle$ denotes the inner product of a and b. In addition, w satisfies the condition that its length is the same as the length from v_d to the plane; this condition can be written as $\langle w, v_d - v_0 \rangle = \langle w, w \rangle$. Let u be the vector that is parallel to w but satisfying $\langle u, v_d - v_0 \rangle = 1$. Thus, u is the solution to the square system of linear equations

$$\begin{bmatrix} v_1 - v_0 \\ v_2 - v_0 \\ \vdots \\ v_{d-1} - v_0 \\ v_d - v_0 \end{bmatrix} u = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix} \tag{2}$$

where u is regarded as a column vector.

We would now like to obtain an upper bound on $\|u\|$. Let A denote the $d \times d$ coefficient matrix in (2). Note that (2) implies that u is the last column of A^{-1}. Therefore, $\|u\|$ is equal to the (d, d) entry of $A^{-T}A^{-1}$, that is, the inverse of AA^T. The (d, d) entry of the inverse of AA^T may be obtained by the adjoint formula for the inverse:

$$\|u\|^2 = \det(\bar{A}\bar{A}^T)/\det(AA^T) \tag{3}$$

where $\bar{A}$ denotes the first $d-1$ rows of A. We obtain an upper bound on the numerator by forming the product of its diagonal entries (recall that the determinant of a symmetric positive definite matrix is bounded above by the product of its diagonal entries). The ith diagonal entry of AA^T is the square of the length of the edge from v_i to v_0. Thus, $\det(\bar{A}\bar{A}^T) \le s^{2(d-1)}$. Next, we need a lower bound on the denominator of (3), which is equal to $\det(A)^2$. We see that $\Delta = |\det(A)|$ by (1). Thus, the denominator of (3) is Δ^2.

Combining all this, we conclude that $\|u\| \le s^{d-1}/\Delta$. Recall that w is a multiple of u scaled so that $\langle w, v_d - v_0 \rangle = \langle w, w \rangle$. Because u satisfies $\langle u, v_d - v_0 \rangle = 1$, it is easily seen that we must take $w = u/\|u\|^2$. Thus, $\|w\| = 1/\|u\|$, hence $\|w\| \ge \Delta/s^{d-1}$. ∎

5 Bounds on the coordinate entries of close points

In this section we begin the analysis of the aspect ratio. The purpose of this section is to obtain bounds on the coordinate entries of the vertices that occur in a non-null chain. These bounds are used in later sections to obtain constraints that the ϵ_i's must satisfy. The reader may wish to consult Figure 3, Figure 4 and Figure 5 for informal examples of these constraints before reading these sections.

The main idea of the bounds in this section can be summarized as follows. Recall that a chain is a sequence of nested box faces. Therefore, a face B_i in the chain extends over one additional dimension, say dimension j, compared to B_{i-1}. In this section we show that the close point for B_i must have a jth coordinate entry that is a sizeable distance from B_{i-1}, else the simplex would be null.

Let us introduce notation for this section and the upcoming sections. Suppose that $B_0 \subset B_1 \subset \cdots \subset B_d$ form a chain of box faces that determine non-null simplex, i.e., a

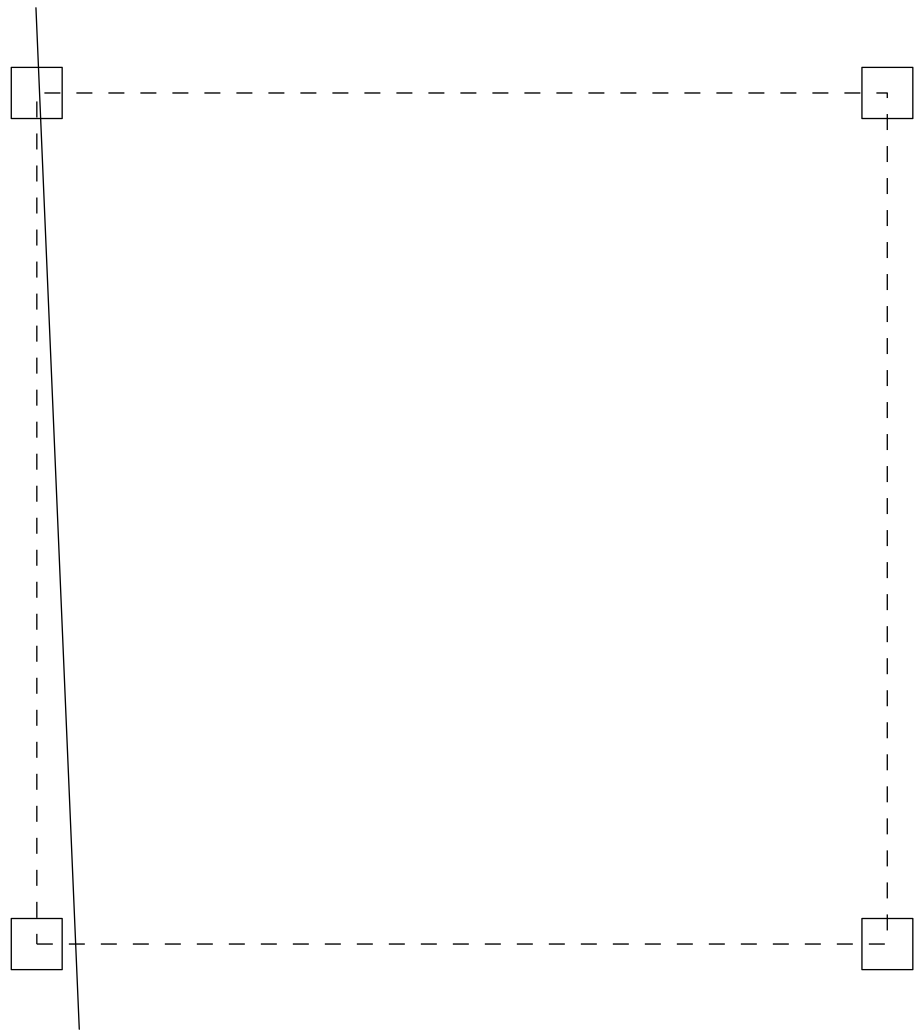

Figure 3: This figure shows what goes wrong if ϵ_0 is too small. In this figure, $d = 2$ and $k = 1$, the edges of the square are dashed, and the ϵ_0-neighborhoods of the vertices are the small squares. If ϵ_0 is very small as in the figure, then we get a triangle with bad aspect ratio, namely, the skinny triangle chopped off the left edge of the square. Thus, we want ϵ_i for each $i = 0, \ldots, d - k - 1$ to be as large as possible subject to other constraints.

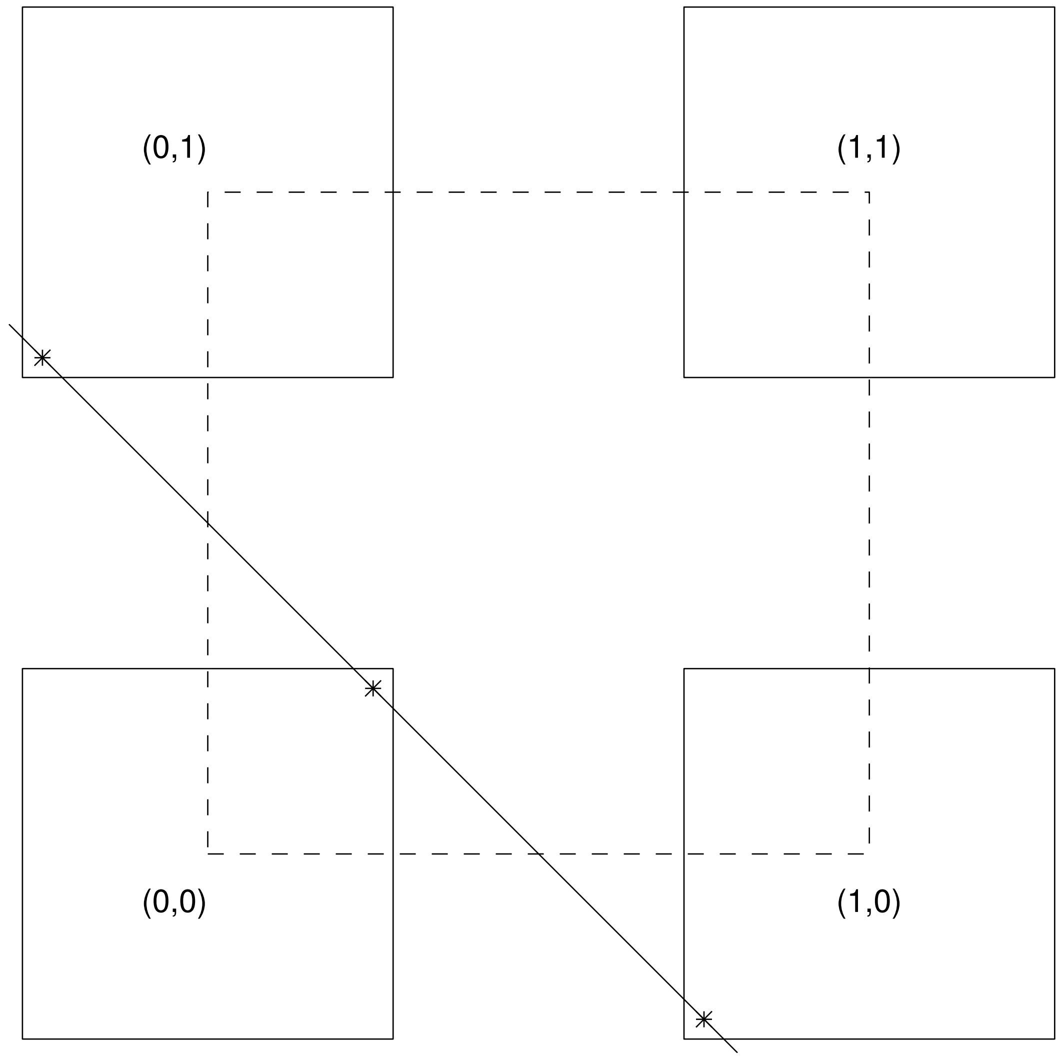

Figure 4: This figure shows what goes wrong if ϵ_0 is too large. In the above example of the $d = 2, k = 1$ case, if ϵ_0 is too large, then we can get a chain with three different close points (marked with asterisks) lying on F. For instance, in the above figure we see that F has a close point for vertex $(0,0)$, a close point for vertex $(1,0)$ which is also the close point for edge $[0,1] \times \{0\}$ (assuming a certain priority rule), and finally a close point for vertex $(0,1)$ which is the close point for the whole square. Thus, we would try to form a simplex that is the convex hull of these three distinct close points. This simplex would have infinite aspect ratio because the points are colinear. Notice that this simplex is not "null"; it does not have a repeated vertex. Therefore, it cannot be dropped from the triangulation because the result would no longer be a simplicial complex. It is clear from the figure that we can prevent this difficulty by making ϵ_0 sufficiently small with respect $1 - \epsilon_0$.

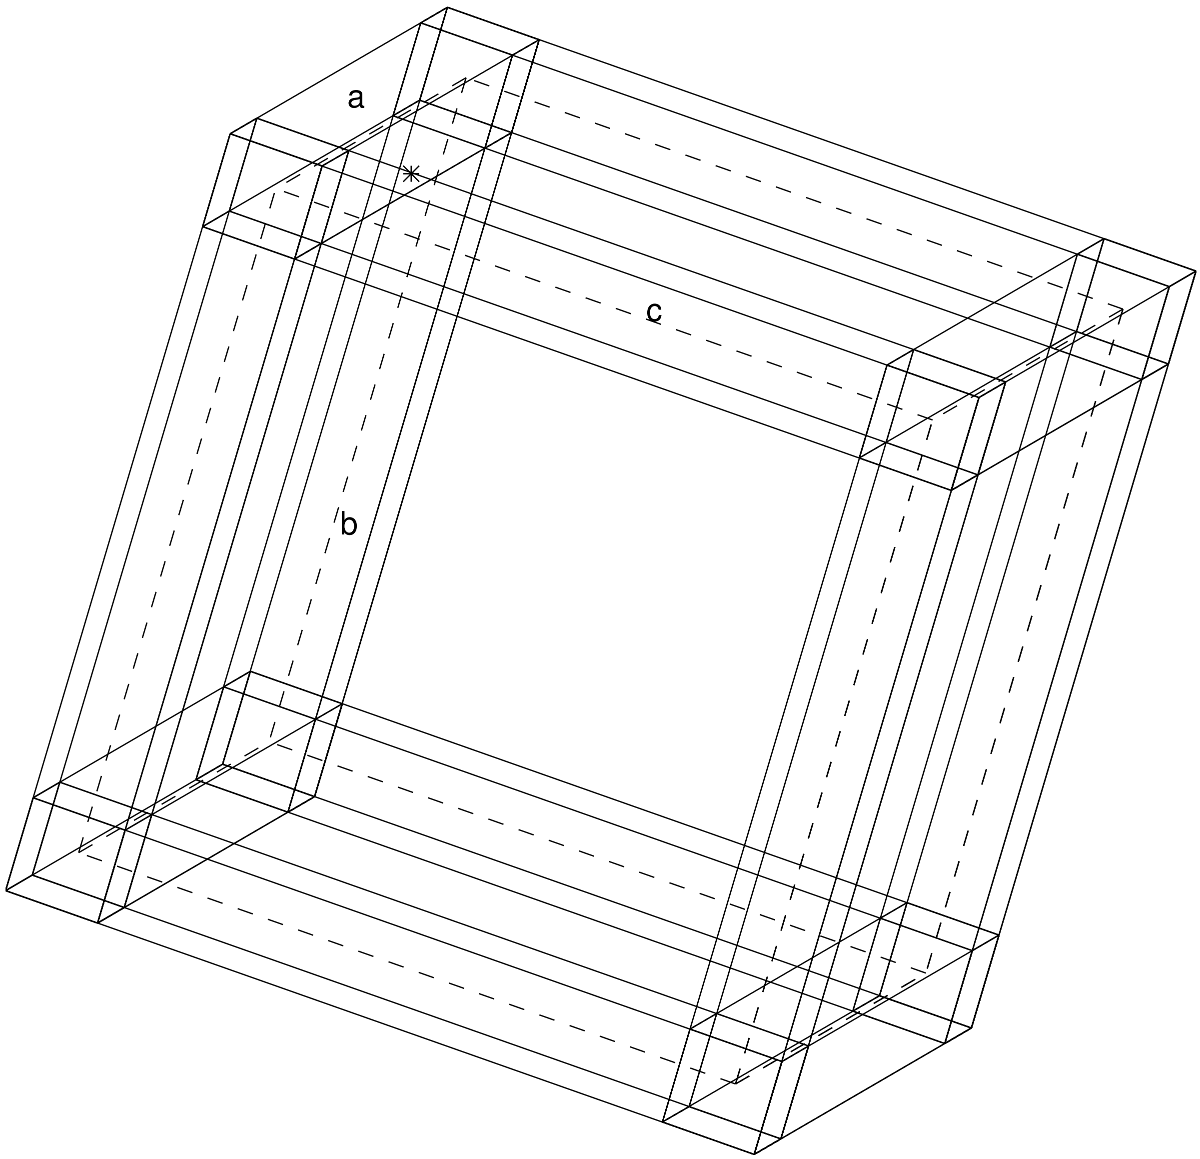

Figure 5: This figure shows what goes wrong if ϵ_1 is too large with respect to ϵ_0 in the $d = 3, k = 1$ case. The figure shows a planar projection of the 3-cube. The neighborhoods around vertices and edges are drawn with $\epsilon_1 = \epsilon_0$. The set F, a line, is perpendicular to the plane of the page at the point marked with an asterisk. Notice that F meets the neighborhoods of three edges of cube marked a, b, c but misses the neighborhood of every vertex. This results in a chain whose associated simplex has infinite aspect ratio as in Figure 4. In order to prevent this, we must make ϵ_0 sufficiently larger than ϵ_1.

simplex with no repeated vertices. By reflecting, translating, and permuting indices, we can assume without loss of generality (to simplify notation) that the chain is

$$\begin{aligned}
B_0 &= \{0\} \times \{0\} \times \cdots \times \{0\}, \\
B_1 &= [0,1] \times \{0\} \times \cdots \times \{0\}, \\
B_2 &= [0,1] \times [0,1] \times \{0\} \times \cdots \times \{0\}, \\
&\vdots \\
B_d &= [0,1] \times \cdots \times [0,1].
\end{aligned}$$

For $i = 0, \ldots, d$, let v_i be $\overline{\mathrm{cl}}(B_i)$. The assumption that the simplex is non-null means that the v_i are all distinct. Let l be the lowest index such that B_l has a close point. Then all of $B_l, \cdots, B_d$ have a close point by the inductiveness of (R1). (Let $l = d + 1$ if no box in the chain has a close point.)

Assume the v_i's are written as row vectors, and let V be the $(d + 1) \times d$ matrix

$$V = \begin{bmatrix} v_0 \\ \vdots \\ v_d \end{bmatrix}.$$

Let the rows of V be numbered 0 to d, and the columns 1 to d. The purpose of this section is to derive some bounds on the entries of V.

We start by looking at the top l rows of V. Note that each of these rows is a box vertex, hence all the entries in these rows are 0's and 1's. Consider row i of V for $i = 0, \ldots, l - 1$ (assuming $l > 0$). Row i is a box vertex lying on B_i. For any $j = i + 1, \ldots, d$, $V_{ij} = 0$ because all points in B_i have zeros in positions $i + 1, \ldots, d$.

We claim further that for $i = 1, \ldots, l - 1$, $V_{ii} = 1$. Suppose not; suppose $V_{ii} = 0$. This means that v_i is a vertex in B_{i-1} as well as in B_i. But recall that v_i is the highest priority vertex in B_i. Since B_i is a superset of B_{i-1}, we conclude that v_i must also be the highest priority vertex in B_{i-1}. This in turn implies that $v_i = v_{i-1}$, contrary to our assumption that the simplex is non-null. This proves that $V_{ii} = 1$. This concludes the analysis of the first l rows of V. For example, if $l = d + 1$ then V has the form

$$V = \begin{bmatrix} 0 & 0 & \cdots & \cdots & 0 \\ 1 & 0 & & & 0 \\ \times & 1 & 0 & & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \times & & \times & 1 & 0 \\ \times & \cdots & \cdots & \times & 1 \end{bmatrix}$$

where each '$\times$' is either 0 or 1.

Now we consider rows of V numbered l to d. Recall that each of these rows has a close point. This close point was determined by either (R1) or (R2) in each case. Let $C_l, \ldots, C_d$ be subfaces of $B_l, \ldots, B_d$ defined as follows. If (R1) was applied to B_i, then let C_i be the high-priority proper subface of B_i that determines v_i. If (R2) was applied to B_i, then let $C_i := B_i$. Thus, in either case, v_i is the close point of C_i, and $v_i \in N(C_i)$.

Let the dimension of C_i, for $i = l, \ldots, d$ be denoted as $r(i)$. Note that $r(i) \le d - k$ for any i as proved in Section 3. We claim that $r(l) \ge r(l+1) \ge \cdots \ge r(d)$. The reason is as follows. Consider an $i \in \{l, \ldots, d-1\}$. Notice that C_i, as a subface of B_i, is also a subface of B_{i+1}. Therefore, C_{i+1} must have higher or equal priority compared to C_i since C_{i+1} was chosen as the high-priority subface of B_{i+1}. Thus, C_{i+1} must have lower or equal dimension than C_i because lower dimensional faces have higher priority.

Now, fix a particular i between l and d. We claim that $|V_{ij}| \le \epsilon_{r(i)}$ for each $j = i+1, \ldots, d$. This because v_i lies in $N(C_i)$, and C_i is a subface of B_i. Therefore, all points of C_i are zero in positions $i+1, \ldots, d$. Thus, a point in $N(C_i)$ has absolute value at most $\epsilon_{r(i)}$ in these positions.

Now consider the diagonal entry V_{ii} for some $i > 0$ between l and d. There are two cases; either $r(i) = 0$ or $r(i) > 0$. If $r(i) = 0$ then C_i is a box vertex. By the same argument as above, this vertex must have a '1' in the ith position (otherwise v_i would also be the close point for B_{i-1}). Therefore, the ith entry of v_i is at least $1 - \epsilon_0$.

The other case is $r(i) > 0$. In this case, we claim that $V_{ii} > \epsilon_{r(i)-1}$. There are three subcases: in the ith coordinate, either C_i is identically 0, identically 1, or extends over dimension i.

If C_i is identically 0 in the ith coordinate, then C_i is actually a subface of B_{i-1}. But then, as above, it must be the highest priority subface of B_{i-1}, contradicting the assumption that the simplex is not null. So this case is impossible.

If C_i is identically 1 in the ith coordinate, then every point $x \in N(C_i)$ has $x_i \ge 1 - \epsilon_{r(i)} > 0.5$ because all of the ϵ_i's are less than 0.5. Thus, $V_{ii} > 0.5 > \epsilon_{r(i)-1}$.

If C_i extends over the ith coordinate, then it has a proper subface C_i' obtained by fixing its ith coordinate to zero. The dimension of C_i' is $r(i) - 1$. Notice that v_i cannot lie in $N(C_i')$ because C_i' has higher priority than C_i and is also a subface of B_i. But notice that if V_{ii} were $\epsilon_{r(i)-1}$ or smaller, then indeed v_i would lie in $N(C_i')$ because $\epsilon_{r(i)-1} \ge \epsilon_{r(i)}$. Therefore, to avoid this contradiction, we must have $V_{ii} > \epsilon_{r(i)-1}$.

To unify this case with the case of $r(i) = 0$, let us adopt the notational convention that $\epsilon_{-1} = 1 - \epsilon_0$. Then we can say that in all cases, $V_{ii} \ge \epsilon_{r(i)-1}$ for $i = \max(l, 1), \ldots, d$.

Finally, we would like to estimate entries of V below the diagonal. The crude upper bound that

$$|V_{ij}| \le 1 + \epsilon_0$$

follows because all the v_i's lie within ϵ_0 of the unit cube.

This conclude our estimates of the entries of V. To summarize:

- The "superdiagonal" entries, that is, entries V_{ij} with $j > i$, satisfy:
  - $V_{ij} = 0$ if $i < l$;
  - $|V_{ij}| < \epsilon_{r(i)}$ if $i \ge l$.
- The "diagonal" entries, that is, entries V_{ii} for $i \ge 1$, satisfy:
  - $V_{ii} = 1$ if $i < l$;
  - $V_{ii} \ge \epsilon_{r(i)-1}$ if $i \ge l$.

- The "subdiagonal" entries, that is, entries V_{ij} with $j < i$, satisfy $|V_{ij}| \leq 1 + \epsilon_0$.

Finally, we have showed that $r(l) \geq r(l+1) \geq \cdots \geq r(d)$.

6 A bound on the number of close points

In this section we show that $l \geq d - k$, i.e., at most $k+1$ of the boxes in the chain can have close points. To prove this requires bounds on the ϵ_i's which are derived below. (Recall the example of Figure 4, in which $d = 2, k = 1$, and $l = 0$. The result of this section is that $l \geq 1$ must hold for the $d = 2, k = 1$ case provided ϵ_0 is chosen correctly.)

Lemma 4 *Suppose the sequence of ϵ_i's satisfies* (12) *below. Then at most $k+1$ box faces in the chain can have close points (assuming, as in Section 5, that no close points in the chain are repeated).*

Proof. Suppose not; suppose that the last $k+2$ boxes, that is, boxes $B_{d-k-1}, \ldots, B_d$ all have close points $v_{d-k-1}, \ldots, v_d$. Consider forming the $(k+1) \times d$ matrix

$$\bar{Y} = \begin{bmatrix} v_{d-k} - v_{d-k-1} \\ v_{d-k+1} - v_{d-k-1} \\ \vdots \\ v_d - v_{d-k-1} \end{bmatrix}.$$

Notice that each row of $\bar{Y}$ is the difference between two points lying on F; therefore each row lies in a k-dimensional vector space. Thus, the maximum possible rank of $\bar{Y}$ is k. Define Y to be the submatrix formed by the last $k+1$ columns of $\bar{Y}$. Thus, Y is a $(k+1) \times (k+1)$ matrix whose rank is also at most k. This means that the determinant of Y is zero. We will show that for appropriate choices of the ϵ_i's, it is impossible for this determinant to be zero, yielding a contradiction.

Let us number both the rows and columns of Y with the index sequence $d-k, \ldots, d$ since this is the natural index sequence inherited from V. Recall that the determinant of Y is a sum of $(k+1)!$ terms. We show that the term from the diagonal must be positive and sufficiently large that the combination of all the other terms could not cancel it out.

Let this diagonal term be denoted as θ:

$$\theta = Y_{d-k,d-k} \cdots Y_{d,d}. \tag{4}$$

Recall that $Y_{ii} = V_{ii} - V_{d-k-1,i}$. Therefore, using the bounds from the last section, we conclude that

$$Y_{ii} \geq \epsilon_{r(i)-1} - \epsilon_{r(d-k-1)}. \tag{5}$$

Let t be some other term, not the diagonal term, in the determinant of Y. Thus,

$$t = Y_{d-k,\pi(d-k)} \cdots Y_{d,\pi(d)} \tag{6}$$

where π is a permutation of $\{d-k, \ldots, d\}$. Let j be the first index such that $\pi(j) \neq j$, so that j lies between $d-k$ and d. (Such a j must exist because t is not the diagonal term).

Note that for this particular j, $\pi(j) > j$, because if $\pi(j) < j$ then the permutation would repeat a column index from a previous factor. Thus, by the bounds in the previous section, we know that

$$|Y_{j,\pi(j)}| \leq \epsilon_{r(j)} + \epsilon_{r(d-k-1)}. \tag{7}$$

For $j' > j$, we get the crude upper bound

$$|Y_{j',\pi(j')}| \leq 1 + 2\epsilon_0 \tag{8}$$

because $v_{j'}$ and v_{d-k-1} both lie in the rectangle $[-\epsilon_0, 1+\epsilon_0]^d$.

Combining (4) and (5) yields

$$\theta \geq Y_{d-k,d-k} \cdots Y_{j-1,j-1} \cdot (\epsilon_{r(j)-1} - \epsilon_{r(d-k-1)}) \cdots (\epsilon_{r(d)-1} - \epsilon_{r(d-k-1)}).$$

Since the r's are in monotone nonincreasing order, we know that $r(d-k-1) \geq r(j) \geq r(j+1) \geq \cdots \geq r(d)$. Thus, each factor in the preceding inequality from the jth onward is greater than or equal to $\epsilon_{r(j)-1} - \epsilon_{r(j)}$. Thus, the preceding inequality implies

$$\theta \geq Y_{d-k,d-k} \cdots Y_{j-1,j-1} \cdot (\epsilon_{r(j)-1} - \epsilon_{r(j)})^{d-j+1}. \tag{9}$$

On the other hand, if we substitute (7) and (8) into (6), we obtain:

$$|t| \leq Y_{d-k,d-k} \cdots Y_{j-1,j-1} \cdot (2\epsilon_{r(j)}) \cdot (1+2\epsilon_0)^{d-j}. \tag{10}$$

Combining (9) and (10), we conclude that we can be assured that

$$\theta > |t| \cdot ((k+1)! - 1) \tag{11}$$

provided that

$$2\epsilon_{r(j)} \cdot (1+2\epsilon_0)^{d-j} \cdot ((k+1)! - 1) < (\epsilon_{r(j)-1} - \epsilon_{r(j)})^{d-j+1} \tag{12}$$

If (11) holds, then the sum of the other $(k+1)! - 1$ terms of the determinant cannot cancel out θ and the determinant of Y cannot be zero. This in turn implies that B_{d-k-1} cannot have a close point, i.e., $l \geq d-k$. ∎

7 Choosing the ϵ_i's

We now explain how to choose the ϵ_i's to satisfy (12). First, we want to get rid of j as a free index. Recall that $j \geq d-k$. Thus $d-j \leq k$. Therefore, if the following bound holds for each $r = 0, \ldots, d-k-1$:

$$2\epsilon_r \cdot (1+2\epsilon_0)^k \cdot ((k+1)! - 1) < (\epsilon_{r-1} - \epsilon_r)^{k+1} \tag{13}$$

this would imply (12). Furthermore, by making the inequality stricter, we can get a positive lower bound on $\det Y$. For example, we could require more strongly that

$$2\epsilon_r \cdot (1+2\epsilon_0)^k \cdot (k+1)! \leq (\epsilon_{r-1} - \epsilon_r)^{k+1} \tag{14}$$

which would give a lower bound of $\theta/(k+1)!$ on $\det Y$.

To define the ϵ_r's, assume d and k are fixed. We now solve (14) with $r = 0$ to obtain ϵ_0 (recalling that ϵ_{-1} is taken to be $1 - \epsilon_0$). Next, we would find ϵ_1 from ϵ_0, and so on, up to ϵ_{d-k-1}.

It does not seem possible to solve (14) in closed form to get an equation for ϵ_r in terms of ϵ_{r-1}. We can get some crude estimates for general d, k, r as follows. Let us assume first that $\epsilon_r \leq 0.5\epsilon_{r-1}$ for each r. Then the right-hand side of (14) has a lower bound of $(0.5\epsilon_{r-1})^{k+1}$. Thus, (14) is satisfied by the recurrence:

$$\epsilon_r := \frac{(0.5\epsilon_{r-1})^{k+1}}{2 \cdot (1+2\epsilon_0)^k \cdot (k+1)!}.$$

To make further estimates, let us overestimate $\epsilon_0 = 0.5$ in the factor $(1+2\epsilon_0)^k$:

$$\begin{aligned} \epsilon_r &:= \frac{(0.5\epsilon_{r-1})^{k+1}}{2 \cdot 2^k \cdot (k+1)!} \\ &= \frac{\epsilon_{r-1}^{k+1}}{2^{2k+2} \cdot (k+1)!}. \end{aligned}$$

To start this recursion off with $r = 0$, we can estimate ϵ_{-1} as 0.5 (recall that $\epsilon_{-1} = 1 - \epsilon_0$; a suitable lower bound on ϵ_{-1} would be 0.5 since $\epsilon_0 < 0.5$). Then the recurrence can be exactly solved starting with $\epsilon_{-1} = 0.5$ to obtain

$$\epsilon_r = 0.5^{(3k+3)(k+1)^r + (2k+2)(k+1)^{r-1} + \cdots + (2k+2)(k+1)^0} \cdot [1/(k+1)!]^{(k+1)^r + (k+1)^{r-1} + \cdots + (k+1)^0}.$$

The maximum value of r is $d - k - 1$ and the maximum value of $k + 1$ is d. Thus, we obtain as a worst-case estimate $\epsilon_r \geq 2^{-2^{O(d \log d)}}$. This bound is double-exponentially small in d because the second-level exponent is polynomial in d.

Another approach to (13) is to solve numerically for particular values of d and k. We find a large value of ϵ_0 satisfying (13) using the fact that $\epsilon_{-1} = 1 - \epsilon_0$. Then we find a large value of ϵ_1, and so on. We have carried out this computation for $d = 3$ and arrived at the following values that satisfy (13):

For the case $d = 3, k = 1$, $\epsilon_0 = 0.15$, $\epsilon_1 = 0.007$, $\epsilon_2 = \epsilon_3 = 0$.

For the case $d = 3, k = 2$, $\epsilon_0 = 0.05$, $\epsilon_1 = \epsilon_2 = \epsilon_3 = 0$.

8 An aspect ratio bound

Now, finally, we can combine the ingredients to obtain a bound on the aspect ratio of any simplex in our triangulation. According to Lemma 3, it suffices to obtain a positive lower bound on the determinant of the matrix

$$W = \begin{bmatrix} v_1 - v_0 \\ v_2 - v_0 \\ \vdots \\ v_d - v_0 \end{bmatrix}.$$

This matrix can be simplified because we have now proved that $v_0 = (0, \ldots, 0)$. This follows from Lemma 4: since $d - k \geq 1$, B_0 does not have a close point. Thus, we need a lower bound on the determinant of

$$W = \begin{bmatrix} v_1 \\ v_2 \\ \vdots \\ v_d \end{bmatrix}.$$

Following the analysis of Section 5, we conclude that W can be written in the form

$$W = \begin{bmatrix} L & 0 \\ X & R \end{bmatrix}$$

where L is an $(l-1) \times (l-1)$ unit lower triangular matrix. Therefore, $\det(W) = \det(R)$, so it suffices to obtain a positive lower bound on the determinant of the $(d-l+1) \times (d-l+1)$ matrix R.

The worst case bound on the determinant of R comes when $l = d - k$, so that R is a $(k+1) \times (k+1)$ matrix. This matrix R has a very similar structure to the matrix Y analyzed in Section 6. In particular, the (i, i) diagonal entry of R is at least $\epsilon_{r(i)-1}$, whereas the (i, j) superdiagonal entry is at most $\epsilon_{r(i)}$ in magnitude. The subdiagonal entries are at most $1 + \epsilon_0$. These bounds are in fact better than the bounds on Y, in other words, the diagonal entries of R have better lower bounds than the corresponding diagonal entries of Y, and similarly, the off-diagonal entries have better upper bounds. Thus, the determinant of R is at least $1/(k+1)!$ multiplied by the diagonal term, because the same result held for Y if (14) holds. The worst case is when all the diagonal entries of R are ϵ_{d-k-1}. Thus, a lower bound on the determinant of W is $\epsilon_{d-k-1}^{k+1}/(k+1)!$ Recall that ϵ_{d-k-1} is double-exponentially small in d. Therefore, raising this to the $k+1$ power and dividing by $(k+1)!$ yields a lower bound on the determinant that is also double-exponential in d.

9 Open questions

Here are some open question raised by this work:

1. Can the bounds be improved? Is the aspect ratio bound inherently double exponential in d (i.e. are there lower bounds)? How much can the actual numerical bounds be improved for $d = 3$?

2. Can this work be generalized to the case that F is a curved manifold? This would be a first step towards generalizing the QMG algorithm to handle curved surfaces. In the curved case, it is presumably necessary that the grid of boxes be sufficiently small that F bends only a small amount in each box.

3. Is it possible to triangulate a grid of boxes cut by a hyperplane in such a way that all the dihedral angles of all simplices are bounded by $\pi/2$? Dihedral bounds of $\pi/2$ are important for some applications (see e.g. [13]) but it is not known how to triangulate

even the special case of convex three-dimensional polyhedra and obtain such an angle bound. (The latter problem was solved in two dimensions by [1]; see also [2], [4] and [8].)

References

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

[2] M. Bern, P. Chew, D. Eppstein, and J. Ruppert. Dihedral bounds for mesh generation in high dimensions. In *Proc. 6th ACM-SIAM Symposium on Discrete Algorithms*, pages 189–196, 1995.

[3] M. Bern, D. Eppstein, and J. Gilbert. Provably good mesh generation. In *Proc. 31st IEEE Symp. Foundations of Computer Science*, pages 231–241, 1990.

[4] M. Bern, S. Mitchell, and J. Ruppert. Linear-size nonobtuse triangulations of polygons. In *Proc. 10th ACM Symp. on Computational Geometry*, pages 221–230, 1994.

[5] L. P. Chew. Guaranteed-quality triangular meshes. Technical Report 89–983, Department of Computer Science, Cornell University, Ithaca, New York, 1989.

[6] P. Frey, B. Sarter, and M. Gautherie. Fully automatic mesh generation for 3–D domains based upon voxel sets. *International Journal of Numerical Methods in Engineering*, 37:2735–2754, 1994.

[7] C. Johnson. *Numerical solution of partial differential equations by the finite element method.* Cambridge University Press, 1987.

[8] E. Melissaratos and D. Souvaine. Coping with inconsistencies: A new approach to produce quality triangulations of polygonal domains with holes. In *Proc. 8th ACM Symp. on Computational Geometry*, pages 202–211, 1992.

[9] S. A. Mitchell and S. A. Vavasis. Quality mesh generation in three dimensions. In *Proc. 8th ACM Symp. on Computational Geometry*, pages 212–221, 1992.

[10] S. A. Mitchell and S. A. Vavasis. Quality mesh generation in higher dimensions. In preparation, 1995.

[11] M. J. Todd. *The Computation of Fixed Points and Applications.* Springer–Verlag, Berlin, 1976.

[12] S. Vavasis. QMG: a finite element mesh generation package. Available at `http://www.cs.cornell.edu /Info /People /vavasis /qmg-home.html`, 1995.

[13] 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.*