

4-1-2005

Lower Bounds for Simplicial Covers and Triangulations of Cubes

Adam Bliss '03
Harvey Mudd College

Francis E. Su
Harvey Mudd College

Recommended Citation

Adam Bliss and Francis Edward Su. Lower bounds for simplicial covers and triangulations of cubes. *Discrete Comput. Geom.*, 33(4):669–686, 2005.

This Article - preprint is brought to you for free and open access by the HMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in All HMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

LOWER BOUNDS FOR SIMPLICIAL COVERS AND TRIANGULATIONS OF CUBES

ADAM BLISS * AND FRANCIS EDWARD SU **

ABSTRACT. We show that the size of a minimal simplicial cover of a polytope P is a lower bound for the size of a minimal triangulation of P , including ones with extra vertices. We then use this fact to study minimal triangulations of cubes, and we improve lower bounds for covers and triangulations in dimensions 4 through at least 12 (and possibly more dimensions as well). Important ingredients are an analysis of the number of exterior faces that a simplex in the cube can have of a specified dimension and volume, and a characterization of corner simplices in terms of their exterior faces.

Let P be a convex polytope. A (*simplicial*) *cover* of P is a collection of simplices such that (i) the vertices of the simplices are vertices of P and (ii) the union of the simplices is P . For example, some covers are triangulations in which simplices meet face-to-face and have disjoint interiors; although in general, cover elements may overlap.

Define the *covering number* $C(P)$ to be the minimal number of simplices needed for a cover of a polytope P . Although the covering number is of interest in its own right (see [3]), we prove in Theorem 1 that the covering number of P also gives a lower bound for the size of a minimal triangulation of P , including triangulations with extra vertices.

We then use the covering number to study the classical problem of determining the size of the minimal triangulation when P is a d -dimensional cube (this is sometimes called the *simplicity* of the cube). Let $I = [0, 1]$, and let I^d denote the d -cube. We define the notion of an *exterior face* of a cube simplex, develop a counting function for exterior faces of prescribed dimension and volume, and establish a recursive bound in Theorem 2 that yields an absolute bound for the cases of interest in Theorem 3. This bound can be improved still further by analyzing separately the corner simplices of the cube; Theorem 4 characterizes corner simplices in terms of the number of exterior faces they have. We use these results to establish new lower bounds for $C(I^d)$, and via Theorem 1, these yield the best known bounds for general triangulations in dimensions up through at least 12 (and possibly more dimensions— see the concluding section). For $d = 4$ through $d = 12$ these numbers are:

$$16, 60, 252, 1143, 5104, 22616, 98183, 520865, 2.9276 \times 10^6$$

In particular, our results show that Mara's minimal triangulation of the 4-cube [7], using 16 simplices, is also a minimal simplicial cover; furthermore, using extra vertices will not produce any smaller triangulations of the 4-cube. See Theorem 5. A comparison of our bounds with bounds for other kinds of decompositions may be found in Table 1.

*Research partially supported by a Beckman Research Grant at Harvey Mudd College.

**Research partially supported by NSF Grant DMS-0301129.

This paper is organized as follows. Section 1 discusses the relationship between minimal simplicial covers and minimal triangulations of arbitrary polytopes P . Section 2 gives background on triangulations of cubes. We develop some terminology in Section 3, and in Section 4, we develop constraints for a linear program for our problem. This involves a counting function F which counts the number of exterior faces that a simplex can have; in order to estimate F , we develop in Sections 5 and 6 some theory regarding the way exterior faces of simplices relate to simplices in the cube. Then, in Section 7 we show how this theory produces a recurrence for the counting function that can be used to get bounds on F , and Section 8 refines the earlier linear program. The final section discusses our results and some open questions.

1. MINIMAL COVERS BOUND TRIANGULATIONS WITH EXTRA VERTICES

When speaking of minimal triangulations of a polytope P , we must be careful to distinguish what kind of decomposition we mean, since there are many such notions in the literature. In this paper, a *triangulation* of a d -polytope P will always mean a decomposition of P into d -simplices with disjoint interiors, such that each pair of simplices intersect in a face common to both or not at all.

In the literature, some authors restrict attention to triangulations whose vertices are required to be vertices of the polytope. We refer to these as *vertex triangulations*. Others do not always require triangulations to meet face-to-face; such decompositions are sometimes called *dissections*. They are covers by simplices that do not have overlapping interiors. A *vertex dissection* is a dissection in which vertices of simplices come from those of the polytope.

We let $D(P), T(P), D^v(P), T^v(P)$ denote, respectively, the size of the smallest possible dissection, triangulation, vertex dissection, and vertex triangulation of P . (We note that our notation differs from that of Smith [12], who uses T to include dissections, and of Hughes-Anderson [6], who use T for vertex triangulations.)

The evident inclusions immediately imply that $D(P) \leq T(P) \leq T^v(P)$ and $D(P) \leq D^v(P) \leq T^v(P)$. Also, $C(P) \leq D^v(P)$ since any vertex dissection is a simplicial cover—this does not hold for general dissections, since in a cover of P , vertices of simplices must be vertices of P . What may therefore be somewhat surprising is the following important inequality relating covers to general triangulations:

Theorem 1. *For any convex polytope P , the covering number $C(P)$ satisfies:*

$$C(P) \leq T(P).$$

Some covers arise as images of simplicial Brouwer self-maps of a triangulated polytope and are related to the polytopal Sperner lemma of DeLoera-Peterson-Su [3]. Such considerations form the basis for a proof of Theorem 1.

Proof. Let K be a triangulation of P . For any such triangulation construct a piecewise linear (PL) map $f_K : P \rightarrow P$ in the following way. If v is a vertex of K , define $f_K(v)$ to be

any vertex of P on the smallest-dimensional P -face that contains v . Then extend this map linearly across each simplex of K . See Figure 1.

Thus f_K is a PL-map from P to P that takes simplices of K to simplices formed by vertices of P , and these images must be a cover of P because this map is a Brouwer map of degree 1 [3, Prop. 3]. Thus there are at least $C(P)$ such simplices in the triangulation K . \square




FIGURE 1. The piecewise linear map f_K . Two of the simplices have been shaded to show how they behave under f_K . Note that the images of simplices in a triangulation can overlap.

Thus the covering number is a lower bound for the size of any triangulation (including ones with extra vertices). We make two remarks about the proof. First, the above assignment of f_K to vertices of P is called a *Sperner labelling* of the vertices, in the sense of [3]. Secondly, this proof does not work for arbitrary dissections, because if the simplices do not meet face-to-face, the resulting PL-map may not be well-defined.

We may summarize the relationships discussed above as

$$(1) \quad C(P), D(P) \leq T(P), D^v(P) \leq T^v(P)$$

which signifies a partial ordering in which there is no known relation between C and D nor between T and D^v , but all other inequalities hold. Very little is known about the strictness of the inequalities above for any class of polytopes P . Constructions in [1] show that in some polytopes, $C(P)$ can be strictly smaller than $T^v(P)$, although it is unknown whether this occurs for cubes.

2. TRIANGULATIONS OF CUBES

We now restrict our attention to cubes. Interest in small triangulations of cubes stems principally from certain simplicial fixed-point algorithms (e.g. see [14]) which run faster when there are fewer simplices. The same considerations govern more recent *fair division procedures* [11, 13] that depend on efficient triangulations.

The d -cube $[0, 1]^d$ has a standard triangulation T_0 of size $d!$ in which each permutation $(x_{s_1}, \dots, x_{s_n})$ of n variables x_1, \dots, x_n is associated with the simplex of points for which $0 \leq$

TABLE 1. Comparison of best-known lower bounds for cubes, and kinds of decompositions to which they apply. Equal signs denote the cases for $d \geq 4$ for which bounds are known to be optimal.

Dimension	Smith [12]	our bounds	Hughes [5],
	C, D, T, D^v, T^v		Hughes-Anderson*[6], Cottle**[2] and Sallee**[10] D^v, T^v
3	5	5	5
4	15	$C, T=16$	$D^v=16, T^v=16^{**}$
5	48	60	61, $T^v=67^*$
6	174	252	$270^*, T^v=308^*$
7	681	1,143	$1,175^*, T^v=1493^*$
8	2,863	5,104	5,522
9	12,811	22,616	26,593
10	60,574	98,183	131,269
11	300,956	520,865	665,272
12	1,564,340	2.9276×10^6	
d	asymptotic bound $\frac{6^{\frac{d}{2}} d!}{2^{(d+1)} \frac{d+1}{2}}$		

$x_{s_1} \leq x_{s_2} \leq x_{s_3} \leq \dots \leq x_{s_n} \leq 1$. From this description, it is easy to see that the simplices meet only in common faces and have no common interior points. This triangulation is maximal among those that use only vertices of the cube.

However, only for $d = 2$ is this triangulation minimal. For $d = 3$, it is easy to check that the T^v -minimal triangulation is of size 5, formed by four corner simplices at non-adjacent vertices, and one fat regular tetrahedron using the other four vertices. It follows from results of Smith [12] that this is minimal for all other kinds of decompositions as well. For $d = 4$, Mara [7] produced a triangulation using only 16 simplices; Cottle [2] and Sallee [10] showed it was minimal for $T^v(I^4)$ and Hughes [5] showed $D^v(I^4) = 16$. In Theorem 5, we show that $C(I^4) = 16$ and therefore $T(I^4) = 16$. It is unknown whether $D(I^4) = 15$ or 16. Hughes-Anderson [6] showed that $T^v(I^5) = 67, T^v(I^6) = 308, T^v(I^7) = 1493$, but for $d \geq 8$ there are no exact results for $T^v(I^d)$. For $d \geq 5$, there are no exact results known for C, D, T, D^v . We improve the best bounds known for C and T in many small dimensions. See Table 1.

To obtain an asymptotic bound for the size of the minimal cover of a d -cube, one might use the following simple idea: the number of simplices in the cover is bounded below by the volume of the cube divided by the volume of the largest possible simplex in the d -cube. This volume is always of the form $V(d)/d!$ for $V(d)$ a positive integer. Determining $V(d)$ is a hard problem related to the *Hadamard determinant problem*, but for small d the values are known. See Table 2.

TABLE 2. Some values of $V(d)$ from [4].

d	0	1	2	3	4	5	6	7	8	9	10	11	12	13
$V(d)$	1	1	1	2	3	5	9	32	56	144	320	1458	3645	9477

The survey [4] also gives some infinite families of answers, and an asymptotic upper bound for $V(d)$:

$$(2) \quad V(d) \leq \frac{(d+1)^{\frac{d+1}{2}}}{2^d}$$

from which one may obtain $C(d) \geq d!2^d(d+1)^{-\frac{d+1}{2}}$. Smith [12] observed that one can improve this technique by considering *hyperbolic volumes* instead of Euclidean volumes, and arrived at the improved bound

$$(3) \quad C(d) \geq \frac{6^{\frac{d}{2}}d!}{2(d+1)^{\frac{d+1}{2}}}.$$

(Although Smith only spoke of triangulations and non-overlapping decompositions of the cube, his method actually also applies to covers as well.) This asymptotic bound remains the best asymptotic bound for $C(d)$ for arbitrary d , but our methods improve the explicit bounds that Smith gives for low dimensions; see Table 1.

However, his results are not fully comparable to ours, since his methods also apply to general dissections. Thus our work gives tighter lower bounds for covers and triangulations in specific dimensions.

In contrast to Smith, we avoid the use of hyperbolic geometry. Instead, we develop a linear program whose optimal solution is a bound for the size of the minimal cover. The linear programming approach was initiated by Sallee [9], Hughes [5], and Hughes-Anderson [6] to study triangulations of the cube. However, the results of Hughes and Hughes-Anderson are not fully comparable to ours, because their methods do not apply to covers.

On the other hand, Sallee's method does apply to covers (although he only spoke of vertex triangulations), and so our Theorem 1 shows that his results hold for general triangulations as well. As reported by Hughes [5], Sallee's method gives these lower bounds for $d = 3$ though $d = 11$: 5, 16, 60, 250, 1117, 4680, 21384, 95064, 502289, although the more recent determination of $V(10)$ would improve the last two bounds to 95708, 516465. Thus our bounds agree with his for $d \leq 5$ and dominate his bounds for $d > 5$.

Upper bounds for minimal triangulations can be obtained by construction. Recent work of Orden-Santos [8] shows that the d -cube can be triangulated with $O(.816^d d!)$ simplices, so there remains a large gap between this asymptotic upper bound and the asymptotic lower bound of Smith. A survey of specific upper bounds in low dimensions (up through 12) may be found in [12].

3. EXTERIOR FACES OF SIMPLICES IN THE CUBE

Hereafter when we refer to a simplex, we will (unless otherwise specified) mean a non-degenerate simplex spanned by vertices of the unit d -cube.

A j -face of a d -simplex is the j -simplex spanned by some $j + 1$ of the simplex's vertices. A j -face of a d -cube is the j -cube spanned by some 2^j of the cube's vertices that lie in a j -dimensional hyperplane on which $d - j$ coordinates agree. In both cases the number $d - j$ is called the *codimension* of the face. A face of codimension 1 is called a *facet*. Two j -faces of a d -cube are said to be *parallel* if the j -dimensional hyperplanes containing them are parallel.

We say a j -face of a d -simplex is *exterior* if it is contained in a j -face of the d -cube. The empty set will also be considered an exterior face. As an example, the diagonal of a facet of the 3-cube is *not* an exterior 1-face, because it is not contained in a 1-face of the d -cube.

We can represent a d -simplex σ in a d -cube as a $(d + 1) \times d$ matrix M in which the rows are coordinates for the vertices of σ . We shall call M the *matrix representation* of σ . Let $[1|M]$ denote $(d + 1) \times (d + 1)$ square matrix formed by augmenting M by an initial column of ones. Then $|\det[1|M]|/d!$ is the volume of σ . In particular, since the vertices of σ are chosen from $\{0, 1\}^d$, every entry in M is either a zero or a one, so volumes of d -simplices are always integer multiples of $(1/d!)$. For convenience we shall call this integer $|\det[1|M]|$ the *class* of σ ; it is a kind of normalized volume. Simplices of class 0 are degenerate.

In the matrix M , a choice of any $j + 1$ rows corresponds to a j -face τ of σ . We shall call these rows the *face-rows* of τ ; they represent the vertices of τ . We call all other rows of M the *non-face-rows* of τ .

In a dual fashion, a choice of any j columns corresponds to a choice of a j -face F of the d -cube and all the j -faces parallel to it. We call these columns the *cube-face-columns* of F ; they correspond to the coordinates that vary over the face F , and outside of these columns, the coordinates of points on the face F are fixed.

Thus a j -face τ (of a simplex σ) is exterior if and only if there is a choice of some j columns outside of which the face-rows of τ are identical. Then for exterior faces τ , we may speak of these columns as the *cube-face-columns* of τ , and all the other columns are the *non-cube-face-columns* of τ .

As an example, the following matrix M represents a simplex α in the 5-cube:

$$(4) \quad \begin{array}{cccccc} & & \downarrow & \downarrow & \downarrow & & \\ \Rightarrow & 0 & 0 & 1 & 1 & 0 & \leftarrow \\ & & 1 & 0 & 1 & 1 & 0 & \leftarrow \\ & & 0 & 0 & 0 & 1 & 0 & \leftarrow \\ \Rightarrow & 0 & 1 & 1 & 0 & 0 & & \\ \Rightarrow & 0 & 1 & 1 & 1 & 0 & \leftarrow \\ \Rightarrow & 0 & 1 & 1 & 1 & 1 & & \\ & & \uparrow & \uparrow & \uparrow & & \end{array} .$$

The rows and columns marked by single-arrows on the right and top are face-rows and cube-face-columns for some exterior 3-face σ . One can verify this by checking that there are 3 columns and $3 + 1$ rows; and after deleting the cube-face-columns, the face-rows look identical. Another way to say this is in any fixed non-cube-face-column, the entries in the cube face rows must be the same. (For instance, in column 4, $m_{41} = m_{42} = m_{43} = m_{45}$.) Another exterior 3-face τ is represented by the face-rows and cube-face-columns marked on the left and bottom by double-arrows. It contains an exterior 2-face represented by rows $\{4, 5, 6\}$ and columns $\{4, 5\}$. The rows $\{1, 2\}$ represent an exterior edge because there is a cube-face-column (i.e., column 1) outside of which the two rows are identical. The rows $\{2, 3\}$, however, do not represent an exterior edge.

4. CONSTRAINTS ON COVERS

A simplicial cover of the d -cube induces simplicial covers of each of its j -faces (which are again cubes). Note also that those covers consist of *exterior faces* of the simplices in the original cover. Thus, for each dimension j , a natural constraint for a d -cube cover is that its d -simplices must have enough j -dimensional exterior faces to cover the j -faces of the d -cube.

Definition. Let $F(d, c, d', c')$ count the maximal number of dimension d' , class c' exterior faces that any dimension d , class c simplex in the d -cube can have.

Given a cover of the d -cube, let x_c represent the number of simplices of class c in that cover. We wish to minimize $\sum x_c$, the total number of simplices in the cover, subject to some constraints:

$$(5) \quad \sum_{c=1}^{V(d')} \frac{c}{d!} F(d, c, d', c) x_c \geq 2^{d-d'} \binom{d}{d'}, \quad (d' = 1, 2, \dots, d).$$

Recall that $V(d)/d!$ is the volume of the largest d -simplex in the d -cube. Equivalently, $V(d)$ is the class of that largest simplex.

To see how the constraints arise, observe that the right side of (5) counts the total d' -volume of the cube-faces of dimension d' . For each dimension $d' \leq d$, there are $2^{d-d'} \binom{d}{d'}$ d' -faces in the d -cube, each with d' -volume 1.

So there must be enough elements of the cover to cover the d' -volume in each dimension. Notice that in a cover, the d' -cube-faces must be covered by exterior facets of the exterior $(d' + 1)$ -faces that cover the $(d' + 1)$ -cube-faces. Thus for d' -cube-faces, we need only consider elements of the cover that arise in a successive chain of exterior facets of exterior facets, up through the top dimension. Since Proposition 1 will show that any exterior facet of a simplex must have the same class as that simplex, we need only consider exterior faces that have same class as the simplices they lie on. Thus the left side of (5) counts the maximal volume of the exterior dimension d' -faces that could arise from these elements of the cover, since x_c is the number of d -simplices of class c in the cover, $F(d, c, d', c)$ is the maximal

number of exterior d' -simplices that a d -simplex can have (of the same class c), and $c/d'!$ is the volume of such a simplex.

For these constraints to be helpful, we will need the value of $F(d, c, d', c')$, or at least, an upper bound, and this allows us to improve our linear program further later.

5. PROJECTING ALONG AN EXTERIOR FACE

Let α be a d -simplex of class c in the d -cube. Suppose that σ is an exterior d' -face of α . Without loss of generality, we can assume that one of the vertices of σ is at the origin, and that σ is contained in the d -cube-face in which last $d - j$ coordinates are zero.

Consider the matrix representation M of α . We may assume that the first $j + 1$ rows of M are the face-rows of σ and the first row is the origin. By assumption, the last $d - j$ coordinates of these rows are all zero, i.e., the cube-face-columns of σ are the first j columns of M . Thus M has the following form:

$$(6) \quad M = \left[\begin{array}{ccc|ccc} 0 & \cdots & 0 & 0 & \cdots & 0 \\ \hline & A & & & \text{zeroes} & \\ \hline & C & & & B & \end{array} \right].$$

Here A and B are square submatrices of size $j \times j$ and $(d - j) \times (d - j)$ respectively. The non-degeneracy of σ implies that we can add multiples of the rows containing A to zero out the submatrix C . This yields a new matrix M_σ :

$$(7) \quad M_\sigma = \left[\begin{array}{ccc|ccc} 0 & \cdots & 0 & 0 & \cdots & 0 \\ \hline & A & & & \text{zeroes} & \\ \hline & \text{zeroes} & & & B & \end{array} \right].$$

Note that $[1|M]$ and $[1|M_\sigma]$ have the same determinant.

Let σ^\perp denote the $(d - j)$ -simplex spanned by the points corresponding to the zero vector and the last $d - j$ rows of M_σ (containing the submatrix B). Here, σ^\perp can be viewed as the projection of the simplex α into the orthogonal complement of σ . It is clearly an exterior face of the simplex M_σ of dimension $d - d'$.

Let $\pi_\sigma : [0, 1]^d \rightarrow [0, 1]^d$ denote the projection that zeroes out the first j coordinates; it collapses the simplex σ to the origin, and sends all other vertices of α to σ^\perp . We shall call the map π_σ the *projection along the face σ* . From this discussion we can draw some immediate consequences:

Proposition 1. *If a d -simplex α has an exterior d' -face σ , then the class of σ divides the class of α . In particular, if σ has codimension 1, then the class of σ equals the class of α .*

Proof. Note that the class of σ equals $\det(A)$ and $\det(A)$ divides $\det[1|M_\sigma]$, which equals $\det[1|M]$. The second statement follows from noting that the submatrix B of the matrix M that represents α is a 1×1 matrix that must contain a 1 if it is non-degenerate. Hence the determinant of A equals the determinant of the whole matrix. \square

Proposition 2. *Suppose a non-degenerate d -simplex α has an exterior d' -face σ in the cube-face F . Then no cube-face parallel to F contains more than one vertex of α .*

Proof. If a cube-face parallel to F contained two vertices of α , then two rows of M that intersect the submatrix B would agree in the last $(d - j)$ coordinates. After transforming M to M_σ by row operations, these two rows would be identical. Hence $\det[1|M_\sigma] = 0$, contradicting the non-degeneracy of the simplex α . \square

This immediately implies:

Proposition 3. *A non-degenerate simplex α cannot have exterior j -faces in two parallel cube-faces, for $j > 0$. Thus in the matrix representation of α , a given (non-empty) cube-face-column set belongs to at most one exterior face.*

(Note that a class 2 simplex in the 3-cube does have 1-faces in parallel 2-faces of the 3-cube, but these are not exterior faces by our definition.)

Proposition 4. *The projection π_σ is one-to-one on the vertices of α that are not in σ .*

Proof. The fiber of the projection π_σ on any vertex of σ^\perp is a cube-face parallel to σ ; two vertices of α in one fiber would then contradict Proposition 2. \square

6. FOOTPRINTS AND SHADOWS

Now suppose that, in addition to σ , the simplex α has another exterior j -face τ . In order to count how many such τ there are, we examine what happens to τ under π_σ , the projection along the face σ . Since τ is a face of α , the projection $\pi_\sigma(\tau)$ is a subset of σ^\perp .

Definition. Define the *shadow* of τ with respect to σ to be the projection $\pi_\sigma(\tau)$. Define the *footprint* of τ with respect to σ to be the intersection $\tau \cap \sigma$. (Note that the footprint may be empty.)

Consider the following example. Let α be the simplex in the cube of Figure 2. The three dotted edges of α are its exterior 1-faces: σ , τ_1 , and τ_2 . Suppose that the origin is denoted by O . Then the triangle σ^\perp is the projection of the simplex α into the orthogonal complement of the exterior edge σ .

With respect to σ , the footprint of τ_1 is the point F , and the shadow of τ_1 is the edge S_1 . Similarly, with respect to σ , the footprint of τ_2 is the point F , and the shadow of τ_2 is the edge S_2 . The exterior face σ also has a footprint and shadow with respect to itself: the footprint is σ and the shadow is just the point at the origin O .




FIGURE 2. A simplex in the cube, with exterior edges denoted by dotted lines. The origin is denoted by O . With respect to σ , the footprint and shadow of τ_i are F and S_i , respectively.

The following lemma is crucial in establishing facts about footprints and shadows.

Lemma 5. *Let M be the matrix representation of α . Suppose that j is the number of face-rows of M that σ and τ have in common, and k is the number of cube-face-columns that they have in common.*

(i) *If $j > 0$ then $j = k + 1$. Else if $j = 0$, then $k = 0$.*

(ii) *As long as $k \neq 0$, the number of non-face-rows that σ and τ share must be equal to the number of non-cube-face-columns they share.*

As an example, consider the simplex α represented by the matrix M in (4) with exterior faces σ and τ . In this σ and τ share $j = 2$ face-rows (rows 1 and 5) and $k = 1$ cube-face-columns (column 2). Since $k \neq 0$, we expect that the number of non-face-rows they share (none) should equal the number of non-cube-face-columns they share (also none).

Proof. Suppose σ has $s + 1$ face-rows (and s cube-face-columns) and τ has $t + 1$ face-rows (and t cube-face-columns).

Assume first that $j > 0$. Since the vertices corresponding to common face-rows all lie in an affine subspace defined by the common cube-face-columns, it must be that $j \leq k + 1$ (else α would be degenerate). On the other hand, the vertices of σ and τ together have $s + t + 2 - j$ face-rows. Since $j > 0$, these rows are identical outside of $s + t - k$ columns. Thus $s + t + 2 - j$ vertices of α lie in an affine subspace of dimension $s + t - k$, so the non-degeneracy of α means that $s + t + 2 - j \leq s + t - k + 1$, or $j \geq k + 1$. Thus $j = k + 1$.

If $j = 0$, then σ and τ have no face-rows in common. We can assume that face-rows of σ are zero outside its cube-face-columns \mathcal{C}_σ . Let \mathcal{C}_τ denote the cube-face-columns of τ , and consider the columns $\mathcal{C}_\sigma \cup \mathcal{C}_\tau$. Outside of these columns, we zero out the entries in all but one of the face-rows of τ by subtracting one of the face-rows from all the others. This does not affect the determinant of M . Now the $s + t + 2$ face-rows of σ and τ together lie in the affine subspace determined by $\mathcal{C}_\sigma \cup \mathcal{C}_\tau$, since they agree outside these columns. Hence $s + t + 1 \leq (s + t - k) + 1$, or $k \leq 0$. By definition, k cannot be negative, so k must be zero.

For the final assertion, let j' and k' be the number of non-face-rows and non-cube-face-columns that σ and τ share in common. One may check that $j' = d - 2d' + j - 1$ and

$k' = d - 2d' + k$. Then the first assertions show that $j' = k'$ except when $j = 0$ (and in that case $k = 0$ and $k' = j' + 1$). \square

The lemma has two important corollaries:

Corollary 6. *Let σ and τ be two exterior faces of a non-degenerate simplex α in the cube. Then $\sigma \cap \tau$ is an exterior face, i.e., the footprint of τ in σ is exterior.*

Proof. If $j = 0$ then $\sigma \cap \tau$ is empty, and by definition exterior. Otherwise $j > 0$, so by Lemma 5 $j = k + 1$. This means that the intersection $\sigma \cap \tau$ is of dimension $j - 1 = k$. Since it lies in a cube-face of dimension k , it is an exterior face of α . \square

Corollary 7. *The shadow of τ with respect to σ is an exterior face of σ^\perp .*

Proof. If $j = 0$ then the shadow $\pi_\sigma(\tau)$ has $t + 1$ face-rows, so it has dimension t . By Lemma 5, if $j = 0$ then $k = 0$, so that the t cube-face-columns of τ are unchanged by the projection π_σ ; outside of these columns, the face-rows of $\pi_\sigma(\tau)$ are identical. So the shadow of τ is contained in a cube-face of dimension t , as desired.

If $j > 0$, then the shadow of τ has $t + 1 - j$ non-zero face-rows and one zero face-row (since vertices of σ map to 0 by π_σ) and hence the shadow of τ has dimension $t + 1 - j = t - k$. But Lemma 5 certifies that $j = k + 1$, so then the shadow of τ is contained in a cube-face of dimension $t - k$ as well. \square

For our purposes, the most important property of footprints and shadows is the following:

Proposition 8. *Given a simplex α in the cube, fix an exterior d' -face σ . Then any exterior d' -face τ of α has a unique footprint-shadow pair with respect to σ .*

(Recall that this footprint may be empty.)

Proof. By Proposition 4, under π_σ , every vertex of σ^\perp (except for the origin) has a unique pre-image in α . Thus the vertices of τ not in σ can be determined from the shadow of τ . Clearly, the vertices of τ in σ can be determined from the footprint of τ . \square

7. A RECURSION FOR $F(d, c, d', c')$

Given a d -simplex α of class c in the d -cube, we wish to count how many exterior faces it may have of dimension d' , class c' . Call such a face an exterior (d', c') -face of α . If it cannot have any, then $F(d, c, d', c') = 0$. Otherwise, it has at least one, call it σ .

Proposition 8 shows that once the exterior face σ is fixed, any other exterior d' -simplex τ determines a unique footprint-shadow pair. The footprint and shadow are exterior faces (of σ and σ^\perp respectively), their dimensions must add up to d' , and the products of their classes must be c' . (The previous sentence holds for empty footprints if empty sets are considered to have dimension 0 and class 1.) So the maximal number of exterior (d', c') -faces must be bounded above by the number of footprint-shadow pairs, summed over the dimension and class of the footprint.

Theorem 2. *If $c > V(d)$, $c' > V(d')$, $d' > d$, or $c' \not\leq c$, then $F(d, c, d', c') = 0$. Otherwise, F obeys the recursion*

$$(8) \quad F(d, c, d', c') \leq \sum_{\delta=0}^{d'} \sum_{\gamma=1}^{c'} F(d', c', \delta, \gamma) \cdot F(d - d', c/c', d' - \delta, c'/\gamma),$$

where $F(d, c, 0, 1)$ is taken to be 1.

Proof. If $c > V(d)$, $c' > V(d')$, $d' > d$, or $c' \not\leq c$, then the definition of $V(d)$ and Proposition 1 rule out the possibility of any exterior (d', c') -faces. Otherwise, if there is an exterior (d', c') -face σ , then for any exterior (d', c') -face τ consider the footprint and shadow of τ with respect to σ .

Let δ be the dimension of the footprint, and let γ be its class. Then $F(d', c', \delta, \gamma)$ counts the maximal number of ways in which the footprint of τ could be an exterior face of σ . Similarly, because of Proposition 1, $F(d - d', c/c', d' - \delta, c'/\gamma)$ counts the maximal number of ways that the shadow of τ can be an exterior face of σ^\perp . Since footprint-shadow pairs are unique (Proposition 8), summing over all possible footprint dimensions, the number of such pairs that τ could assume is given by the right hand side of Equation (8).

The case where the footprint is empty merits special caution. In this case, we choose the convention that the footprint be 0-dimensional (not (-1) -dimensional, as one might expect). There are two reasons we do this. First, there is no overlap with the footprint-shadow pairs that arise from $\delta = 0$, because if τ has empty footprint and the footprint of τ' is a vertex v , then the shadows of τ and τ' must be different ($\pi_\sigma(v) = 0$, so one shadow contains 0 and the other does not). Secondly, in both cases the dimension of the shadow must be d' .

This convention simplifies the sum (8), removing the need for an extra term for $\delta = -1$. \square

The recursion yields some closed form expressions for $F(d, c, d', c')$ in some special cases. Let $\Delta(c) = \min\{d : V(d) \geq c\}$. Hence the smallest dimension in which a class- c simplex appears is at least $\Delta(c)$, with equality if $c = V(d)$. Table 3 shows some known values of $\Delta(c)$.

TABLE 3. Some values of $\Delta(c)$. Compare Table 2.

c	1	2	3	4	5	9	32	56	144	320	1458	3645	9477
$\Delta(c)$	0	3	4	5	5	6	7	8	9	10	11	12	13

Theorem 3. *If $c' = c$, then*

$$(9) \quad F(d, c, d', c) \leq \binom{d - \Delta(c)}{d' - \Delta(c)}.$$

Proof. We prove the theorem by induction on $k = d - \Delta(c)$. If $k = 0$, then $d = \Delta(c)$. Hence for any $d' < \Delta(c)$, both sides of (9) are 0, because $\Delta(c)$ is a lower bound for the dimensions

in which a class- c simplex appears. If $d' = \Delta(c)$, then the left side of (9) is at most 1, and the right side is 1, confirming the desired inequality.

Now assume the theorem holds whenever $d - \Delta(c) \leq k - 1$. We shall show it holds for $d - \Delta(c) = k$.

If $d' = d$, then it is easy to check that both sides of (9) are 1, as desired, and if $d' < \Delta(c)$, then both sides are 0, as desired. The remaining cases for d' are if $\Delta(c) \leq d' < d$. Using the recurrence of Theorem 2, the only nonzero terms in the sum have $\gamma = c$ since $c' = c$, and $\Delta(c) \leq \delta \leq d'$, so we have:

$$(10) \quad F(d, c, d', c) \leq \sum_{\delta=\Delta(c)}^{d'} F(d', c, \delta, c) \cdot F(d - d', 1, d' - \delta, 1).$$

Then if $d' = \Delta(c)$, this inequality becomes

$$F(d, c, d', c) \leq F(\Delta(c), c, \Delta(c), c) \cdot F(d - \Delta(c), 1, 0, 1).$$

By definition, $F(d - \Delta(c), 1, 0, 1) = 1$ and $F(\Delta(c), c, \Delta(c), c)$ is 1 if a class c simplex can be realized in dimension $\Delta(c)$ and it is 0 otherwise. In any case, their product is less than or equal to $\binom{d - \Delta(c)}{\Delta(c) - \Delta(c)} = 1$, as desired. Otherwise, $\Delta(c) < d' < d$, so the inductive hypothesis can be used on the right side of (10), since $d' - \Delta(c)$ and $d - d' - \Delta(c)$ are strictly less than $d - \Delta(c)$. Hence:

$$F(d, c, d', c) \leq \sum_{\delta=\Delta(c)}^{d'} \binom{d' - \Delta(c)}{\delta - \Delta(c)} \binom{d - d'}{d' - \delta} = \binom{d - \Delta(c)}{d' - \Delta(c)},$$

as desired. □

In particular, when $c' = c = 1$, the bound in Theorem 3 becomes $\binom{d}{d'}$ and it is achieved by the corner simplices of the d -cube, i.e., simplices spanned by one vertex and all its nearest neighbors. In fact, as we show in the following theorem, if $1 < d' < d$, corner simplices are the only simplices for which the bound on $F(d, 1, d', 1)$ is sharp.

Theorem 4. *If $1 < d' < d$ and $c = 1$, then the bound of Theorem 3 on $F(d, 1, d', 1)$ is sharp for corner simplices. In fact, for $1 < d' < d$, a simplex has strictly more than $\frac{d-1}{d} \binom{d}{d'}$ exterior d' -faces if and only if it is a corner simplex. Hence for non-corner simplices, the number of exterior d' -faces is bounded above by $\lfloor \frac{d-1}{d} \binom{d}{d'} \rfloor$.*

Note that the “only if” implication does not hold if $d' = 1$ or if $d' = d$; for example if $d = 3$ and $d' = 1$, then a non-corner simplex can have as many exterior edges as a corner (three), and if $d = d'$, then any simplex is its own exterior d' -face. Note also that the bound $\lfloor \frac{d-1}{d} \binom{d}{d'} \rfloor$ in the final assertion holds even for non-corner simplices of class $c > 1$, although in such cases, the bound of Theorem 3 is better.

Proof. Let α denote a corner simplex. Without loss of generality its matrix representation M can be written as the identity $d \times d$ matrix augmented with a row of zeroes. Any choice of d' columns of the d columns of M specifies a set of cube-face-columns, and one may easily

locate $d' + 1$ rows of M outside of which those cube-face-columns are identical. These are face-rows for some exterior d' -face of α . Thus there is one such face (and by Proposition 3, at most one) for each of the $\binom{d}{d'}$ choices of cube-face-column sets.

For the converse, let M be a matrix representation of α , and let \mathcal{C}_j denote the set of $d - d'$ consecutive columns of M that start at the j -th column (and wrap around if necessary). For instance, if M has $d = 8$ columns and $d' = 3$, then \mathcal{C}_3 would be the column set $\{3, 4, 5, 6, 7\}$, and \mathcal{C}_6 would be column set $\{6, 7, 8, 1, 2\}$. Let \mathcal{R}_i be defined similarly for rows of M *excluding the $(d + 1)$ -st row*, e.g., in the previous example, even though M has $d + 1 = 9$ rows, \mathcal{R}_6 would be row set $\{6, 7, 8, 1, 2\}$.

We first show that if the number of exterior d' -faces of α is greater than $\frac{d-1}{d} \binom{d}{d'}$, then α has a d -cycle of exterior d' -faces, in other words, there must be an ordering of the columns of M such that each of $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_d$ represents non-cube-face-columns for some exterior d' -face.

Let $d^* = d - d'$. There are $d!$ possible orderings of the columns of M , and we claim that at least one of them will remain as a possible ordering for a d -cycle if at most $1/d$ of the possible $\binom{d}{d'}$ column subsets of size d^* are *not* non-cube-face-columns for some exterior face. This is true because if some subset A of $d^* = d - d'$ columns are *not* non-cube-face-columns for some exterior face, then this restriction rules out $(d - d')!d'd$ of the $d!$ orderings of the columns of M (since there are $(d - d')!$ ways to order A , $d'!$ to order the columns not in A , and d ways to place the columns of A adjacent to each other). Thus if we rule out strictly fewer than $d!/(d - d')!d'd = \binom{d}{d'}/d$ exterior faces, there will still remain an ordering that could occur as an ordering of the columns of M , and in such an ordering, every \mathcal{C}_i is a set of non-cube-face-columns for some exterior face.

Now we show that if $\mathcal{C}_1, \dots, \mathcal{C}_d$ all represent non-cube-face-column sets for exterior d' -faces, then α must be corner simplex.

Consider \mathcal{C}_1 , the first d^* columns of M ; by assumption, these are non-cube-face-columns for some exterior d' -simplex σ_1 . We can assume (by re-ordering rows if needed) that its non-face-rows are \mathcal{R}_1 , the first d^* rows of M . The intersection of the non-face-rows and non-cube-face-columns forms a $d^* \times d^*$ block B_1 in the matrix M . In any single non-cube-face-column, all entries not in B_1 must be identical. Note that, by symmetry, toggling all the column elements in a column (exchanging 1's and 0's) does not change the isomorphism class of the simplex, and we can use this operation, if needed, in columns 1 through d^* so that entries not in B_1 are all 0.

Similarly, consider columns in \mathcal{C}_2 (columns 2 through $d^* + 1$); these are non-cube-face-columns corresponding to an exterior d' -simplex σ_2 . Because $d' > 1$, σ_1 and σ_2 share a common cube-face-column (e.g., column d), so Lemma 5(ii) applies: noting that that σ_1 and σ_2 share $d^* - 1$ non-cube-face-columns, they must also share $d^* - 1$ non-face-rows. We can thus reorder the rows of M so that σ_2 occupies rows in \mathcal{R}_2 (rows 2 through $d^* + 1$). Thus the corresponding $d^* \times d^*$ block B_2 of σ_2 intersects the block B_1 of σ_1 in a block along the diagonal, and as before, we may toggle column elements in non-cube-face-columns of σ_2 so that all entries in those columns but outside B_2 are 0.

In the same way, by cycling through the columns, we can infer from Lemma 5(ii) that each pair of neighboring faces in the d -cycle of exterior d' -faces must share $d^* - 1$ non-face-rows (because they share $d^* - 1$ non-cube-face-columns and $d' > 1$). In particular, since there are d such faces in the cycle, there are only d rows of M that are used as non-face-rows of such faces. Thus there is one row of M that is not a non-face-row of any face, in other words, it is a face-row of *every* face in the cycle.

We can reorder the rows of M in such a way that this common face-row is in the $d + 1$ -st row of M , and the exterior face whose non-cube-face-columns are \mathcal{C}_k has non-face-rows \mathcal{R}_k , for all $1 \leq k \leq d$. Then, by toggling columns if needed, each of the corresponding $d^* \times d^*$ blocks along the diagonal now has zeroes in its columns outside it. This forces row $d + 1$ to be all zeroes, and all non-diagonal entries are also all zero. Then all diagonal entries must be 1 because α is non-degenerate. Hence α must be isomorphic to a corner. \square

8. IMPROVING THE LINEAR PROGRAM

We can now improve the linear program in (5) by noting that Theorem 3 shows that all the coefficients $F(d, c, d', c)$ of x_c are equal for values of c between any two values of V , i.e., $V(k - 1) < c \leq V(k)$. Since the objective function also equally weights the variables x_c , there is an optimal solution to the linear program (5) whose support lives entirely on the variables $x_{V(k)}$ for $k = 1, \dots, d'$, e.g., x_1, x_2, x_3, x_5, x_9 , etc. Setting $y_k = x_{V(k)}$ for $k \geq 2$, the program becomes:

$$(11) \quad \min \sum_{k=2}^d y_k \text{ subject to}$$

$$\sum_{k=2}^{d'} \frac{V(k)}{d'^!} F(d, V(k), d', V(k)) y_k \geq 2^{d-d'} \binom{d}{d'}, \quad d' = 1, 2, \dots, d.$$

We can improve this linear program further by considering the presence of corner simplices separately from other possible exterior class-1 d' -simplices. In the program above, replace y_2 (which counts the total number of class-1 simplices) by $y_1 + y_2$, where y_1 denotes the number of corner simplices in a cover, and y_2 now denotes the number of class-1 non-corners. Applying Theorems 3 and 4, and the observation that y_1 must be bounded above by the total number of vertices of the d -cube (since there is at most one corner simplex per vertex), we obtain the program:

$$(12) \quad \min \sum_{k=1}^d y_k \text{ subject to}$$

$$y_1 \leq 2^d,$$

$$\sum_{k=1}^{d'} \frac{V(k)}{d'^!} [\epsilon^{(k, d)} \binom{d - \Delta(V(k))}{d' - \Delta(V(k))}] y_k \geq 2^{d-d'} \binom{d}{d'}, \quad d' = 1, 2, \dots, d,$$

where $\epsilon(k, d) = 1$ unless both $k = 2$ and $1 < d' < d$, in which case $\epsilon(k, d) = \frac{d-1}{d}$. Note that $\Delta(V(k))$ is almost always k , unless $k \leq 2$, and the floor function only comes into play when $\epsilon(k, d) \neq 1$.

9. CONCLUSION

We solved the linear program (12) using the package `lp_solve`, and rounded non-integer optimal values up to integers. This yields the bounds in Table 1.

It should be possible to improve our bounds for covers by using more information about the kinds of simplices that can occur in the cube and how they fit together, similar to what Hughes [5] and Hughes-Anderson [6] did to study D^v, T^v -minimal triangulations. These programs rely on enumerating specific features of configuration classes of simplices in the cube, or on enumerating the classes themselves. In the latter case, the resulting linear program has one variable for each isomorphism class of configurations, which for $d = 6$ already involves thousands of variables. The Hughes approach [5] does not work beyond $d > 11$ and the Hughes-Anderson program [6] becomes intractable for $d > 7$.

By contrast, the number of variables in our program for $C(I^d)$ is just d . Also, our results for $C(I^d)$ using a comparatively small program compares quite well to their lower bounds for $T^v(I^d)$ and does not require knowledge of the configuration classes that occur in specific dimensions (other than the existence of corners, which occur in every dimension). So our program can be solved for much larger values of d than programs that use extra information about configuration classes in specific dimensions.

Our bounds dominate Smith's bound up through at least $d = 12$, and computational evidence (limited by roundoff errors) show that our bounds dominate Smith's bound (3) for covers and general triangulations in dimensions up through at least $d = 27$, using specific values of $V(d)$ from Table 2 and the asymptotic bound for $V(d)$ from (2). (Bounds for higher d could not be computed due to solver overflow errors.) However, we do not believe that our bounds will exceed Smith's bound asymptotically.

Our bounds do give some new observations in dimension 4. It was known [2] that the minimal triangulation of the 4-cube has size 16, for triangulations using only the vertices of the cube. Our bounds and Theorem 1 then give a stronger result:

Theorem 5. *The minimal simplicial cover of the 4-cube has size 16, and is achieved by a triangulation. Consequently, adding vertices to the 4-cube will not reduce the size of the minimal triangulation.*

In higher dimensions, it may be the case that the size of a minimal simplicial cover of a d -cube is strictly smaller than size of the minimal triangulation. (Such a cover may have more symmetry than the minimal triangulation.) The intuition here is that the largest simplices in high-dimensional cubes tend to overlap, but more coverage might be possible by using such large simplices than by using small non-overlapping ones.

Open Problem. Find a d -cube for which the minimal cover is strictly smaller than the minimal T^v -triangulation.

If this is true, it would suggest that there may be a d -cube for which adding extra vertices helps to make the triangulation smaller.

Open Problem. Find a d -cube for which the minimal T^v -triangulation is reduced in size by adding extra vertices.

One may also ask a similar question about dissections, and note from Table 1 that even in dimension 4 a question remains:

Open Problem. Find a d -cube for which the minimal dissection is strictly smaller than the minimal triangulation. Does the minimal dissection of the 4-cube have 15 simplices or 16 simplices?

We have been pleasantly surprised by the richness of the geometry of cubes.

Acknowledgements. The authors are grateful to Günter Ziegler and Francisco Santos for helpful conversations.

REFERENCES

- [1] A. Below, U. Brehm, J. A. De Loera, and J. Richter-Gebert. Minimal simplicial dissections and triangulations of convex 3-polytopes. *Discrete Comput. Geom.*, 24(1):35–48, 2000.
- [2] R. W. Cottle. Minimal triangulation of the 4-cube. *Discrete Math.*, 40(1):25–29, 1982.
- [3] J. A. De Loera, E. Peterson, and F. E. Su. A polytopal generalization of Sperner’s lemma. *J. Combin. Theory Ser. A*, 100(1):1–26, 2002.
- [4] M. Hudelson, V. Klee, and D. Larman. Largest j -simplices in d -cubes: some relatives of the Hadamard maximum determinant problem. In *Proceedings of the Fourth Conference of the International Linear Algebra Society (Rotterdam, 1994)*, volume 241/243, pages 519–598, 1996.
- [5] R. B. Hughes. Lower bounds on cube simplicity. *Discrete Math.*, 133(1-3):123–138, 1994.
- [6] R. B. Hughes and M. R. Anderson. Simplicity of the cube. *Discrete Math.*, 158(1-3):99–150, 1996.
- [7] P. S. Mara. Triangulations for the cube. *J. Combinatorial Theory Ser. A*, 20(2):170–177, 1976.
- [8] D. Orden and F. Santos. Asymptotically efficient triangulations of the d -cube. *arXiv:math.CO/0204157 v2*, 2002.
- [9] J. F. Sallee. A note on minimal triangulations of an n -cube. *Discrete Appl. Math.*, 4(3):211–215, 1982.
- [10] J. F. Sallee. A triangulation of the n -cube. *Discrete Math.*, 40(1):81–86, 1982.
- [11] F. W. Simmons and F. E. Su. Consensus-halving via theorems of Borsuk-Ulam and Tucker. *Math. Social Sci.*, 45(1):15–25, 2003.
- [12] W. D. Smith. A lower bound for the simplicity of the n -cube via hyperbolic volumes. *European J. Combin.*, 21(1):131–137, 2000.
- [13] F. E. Su. Rental harmony: Sperner’s lemma in fair division. *Amer. Math. Monthly*, 106(10):930–942, 1999.
- [14] M. J. Todd. The computation of fixed points and applications. In *Lecture Notes in Economics and Mathematical Systems*, New York, 1976. Springer-Verlag.

DEPARTMENT OF MATHEMATICS, HARVEY MUDD COLLEGE, CLAREMONT, CA 91711, U.S.A.
E-mail address: `abliss@hmc.edu`

DEPARTMENT OF MATHEMATICS, HARVEY MUDD COLLEGE, CLAREMONT, CA 91711, U.S.A.
E-mail address: `su@math.hmc.edu`