Covering Numbers of the Cubes

by

Adam Bliss

Francis Su, Advisor

Advisor: ____

Second Reader: ____

(Hank Krieger)

August 2003

Department of Mathematics

HARVEY MUDD COLLEGE

Abstract

Covering Numbers of the Cubes

by Adam Bliss

August 2003

How many triangles does it take to make a square? The answer is simple: two. This problem has a direct analogue in dimensions three and higher, but the answers are much harder to find. We provide new lower bounds in dimensions 4 through 13, an asymptotic lower bound which is inferior to the best-known bound in high dimensions, and some new ideas which produce good upper bounds in both low and high dimensions.

Table of Contents

List of Tables

Acknowledgments

I would like to thank my advisor, Professor Francis Su, and my reader, Hank Krieger, for their support and tolerance.

Chapter 1

Introduction

1.1 Overview

Everyone knows that it takes two triangles to make a square, but how many triangle-based pyramids does it take to make a cube? Most people, after considering the issue, guess 6; in fact it can be done with only 5.

If we continue the analogy into dimension d, the triangle becomes the *d-simplex*, or a convex hull of $d + 1$ points in general position (the points are known as its *vertices*). The square becomes the *d-cube*, or the iterated Cartesian product of the unit interval, $[0, 1]^d$. The simplices must *cover* the cube, meaning that their union must contain the cube, and their vertices must be chosen from the vertices of the cube (i.e., their coordinates in each dimension must be 0 or 1).

The question this thesis will attempt to answer, then, is *how many simplices does it take to cover the d-cube?* The answer to this question is the so-called *covering number* of the d-cube and will be represented by $C(d)$. Rather than finding the exact number, we will be satisfied with lower and upper bounds, but we'll try to make them as sharp as possible. We will also examine some more general notions and some more specific ones.

In this thesis we present work done over the Summer of 2002 in dimensions 4, 5, 6, 7, and 9, which produced new lower bounds (see Table 1.1 for a comparison with Smith's bounds, which we will revisit below); research from Fall 2002 which produced new lower bounds in dimensions 8, 10, 11, 12, and 13; and new ideas

Table 1.1: Lower bounds for the covering number of the cube.

Dim	Smith [8]	Summer	Fall
3	5	5	5
4	15	16	16
5	48	60	60
6	174	245	250
7	681	1031	1117
8	2863	2616	4680
9	12811		21384
10	60574		88172
11	300956		494547
12	1564340		2681790
d	$\geq \frac{6^{\frac{d}{2}} d!}{2(d+1)^{\frac{d+1}{2}}}$		

from Spring 2003 which produced asymptotic lower and upper bounds.

1.2 Background

Discrete geometers have long been interested in a problem closely related to the search for covering numbers. They want to know how many simplices it takes to *triangulate* the d-cube. A *triangulation* is a special kind of cover in which the intersection of any two simplices is a face of each. Every triangulation is a cover, so the size of the smallest triangulation is an upper bound for the size of the smallest cover. However, in some polytopes the smallest cover is smaller than the smallest triangulation. A *dissection* is a cover whose simplices have mutually disjoint interiors; minimal dissections for the d-cube have also been considered. Very little work, however, has been done on fully general covers. In fact, the only directly applicable result we found in the literature is relevant only accidentally: Smith in [8] is concerned with various kinds of dissections, but his bound (which we will examine later) actually applies to all covers.

The interest in small triangulations stems principally from certain simplicial fixed-point algorithms (see e.g. [9]) which run faster when there are fewer simplices; however the matter is also intrinsically interesting from the standpoint of discrete geometry. General covers are applicable to the generalized Sperner's lemma put forth in [5].

A rough-and-ready upper bound for the triangulation size of the d-cube is $d!$. This is achieved by a simple construction involving permutations on the dimensions of the cube. In fact, the size-6 cover for the 3-cube that people first think of is usually one of these constructions. However, as mentioned above, there is actually a size-5 cover, which has long been known to be the best possible.

Seeing that the $d!$ bound is not tight in dimension 3, it is natural to look for a triangulation in dimension 4 with fewer than 24 simplices. Mara ([6]) found one in 1976 that used only 16; he thought that this was the smallest possible triangulation but couldn't quite prove it. Cottle ([2]) completed the proof in 1982, using geometric and combinatorial arguments.

Here's a very simple argument that leads to an asymptotic lower bound for cover size: Let $V(d)/d!$ be the d-dimensional volume [1] of the largest d-simplex in the d-cube. Since the total volume of the d-cube is 1, any cover must include at least $d!/V(d)$ simplices. Now, it turns out that computing $V(d)$ is not very easy. The latest results in [3] give exact answers up to $d = 13$ (See Table 1.2), some infinite families of answers, and an asymptotic bound from which we get

$$C(d) \geq \frac{2^d d!}{(d+1)^{\frac{d+1}{2}}}. \tag{1.1}$$

[1] Hereafter, we will write "d-volume" or simply "volume" when the dimension is unambiguous.

Table 1.2: Some values of $V(d)$ from [3]

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

In [8], Smith recasts the same argument using *hyperbolic* volumes, arriving at the improved bound

$$C(d) \geq \frac{6^{\frac{d}{2}} d!}{2(d+1)^{\frac{d+1}{2}}}. \tag{1.2}$$

Smith also gave some explicit bounds in low dimensions; these numbers are in Table 1.1. Our first goal in this thesis is to improve on these bounds by taking into account how simplices in the d-cube relate to simplices in the faces of the cube. In the next chapter we will provide some needed background definitions and some results about simplices in cubes.

Chapter 2

Preliminaries

2.1 Some Definitions

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, provided they lie on a j-dimensional hyperplane disjoint from the cube's interior. In either case the number $d - j$ is called the *codimension* of the face. A face with codimension 1 is called a *facet*, and a face of codimension 2 is called a *ridge*.

A j-face of a d-simplex in the d-cube is *exterior* if it is contained in a j-face of the d-cube.

Two j-faces of a d-cube are said to be *parallel* if the j-dimensional hyperplanes containing them are parallel.

We can represent a d-simplex σ in a d-cube as a $(d+1) \times d$ matrix M in which each row vector gives the coordinates for one of σ's vertices. If we augment M with a column of ones to form a square matrix N (called the *matrix representation* of σ), then $|\det(N)|/d!$ will give the volume of σ. In particular, since the vertices of σ are chosen from $\{0,1\}^d$, every entry in N is either a zero or a one, so the numerator of this fraction is an integer; it is called the *class* of σ. Simplices of class 0 (along with those of dimension 0) are called *degenerate* and unless otherwise specified we will henceforth refer only to non-degenerate simplices and faces in cubes.

In this matrix N, a choice of any $j + 1$ rows corresponds to a j-face of σ. This face is exterior if and only if there is a choice of some $j + 1$ columns (including the column of ones) outside of which the chosen rows are identical. Conversely, a choice of some $j+1$ columns (again, including the column of ones) corresponds to a choice of a j-face of the d-cube and all the j-faces parallel to it; a face of σ contained in any one of these faces will correspond to a choice of some $j+1$ rows which are identical outside of the chosen columns. This interpretation allows us to prove several useful lemmas.

2.2 Some Lemmas

Lemma 2.2.1. *Suppose a non-degenerate d-simplex σ has an exterior j-face in the cube-face f. Then no cube-face parallel to f contains more than one vertex of σ.*

Proof. Suppose for the sake of contradiction that a cube-face g were parallel to f and contained two of σ's vertices. Consider the matrix representation N of σ. Let $c_1, \ldots, c_{j+1}$ be the columns corresponding to the cube-faces parallel to f (where c_{j+1} is the augmented column of ones); let $r_1, \ldots, r_{j+1}$ be the rows corresponding to the exterior face of σ in f; and let s_1 and s_2 be the rows corresponding to the vertices of σ in g. We will perform some elementary row operations on N. First, subtract r_{j+1} from each of $r_1, \ldots, r_j$, causing them to have 0 as the entry in every column but $c_1, \ldots, c_j$. Next, subtract s_2 from s_1, causing the same for it. Then $r_1, \ldots, r_j, s_1$ all lie in the j-dimensional subspace determined by the columns $c_1, \ldots, c_j$; they must therefore be linearly dependent. But since σ is non-degenerate, $\det(N) \neq 0$ and the rows of N must be linearly *independent*, so this is a contradiction. □

Corollary 2.2.2. *A non-degenerate simplex cannot have exterior j-faces in two parallel cube-faces.*

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.

Lemma 2.2.3. *Let σ be a d-simplex of class c, and suppose σ has an exterior j-face τ of class k. Then $k|c$ and $c/k \leq V(d-j)$.*

Proof. Without a loss of generality, we may assume that one of the vertices of τ is the origin. Consider, then, the matrix representation N of σ. Arrange the rows and columns so that the first $j+1$ rows are the vertices of τ, and the columns corresponding to the cube-face containing τ are the first j columns along with the last column of ones. Subtract the origin-row from the rest of the vertices of τ, leaving j rows in a j-dimensional subspace. These rows must be independent since σ is non-degenerate; therefore they can be used to reduce the first j entries in all but the first $j+1$ rows to 0. The determinant of the matrix is then the determinant of the minor of the upper-right-hand entry; this minor is block-diagonal. The determinant of the first block is the class of τ, i.e. k. If we add to the second block a row of zeroes and then a column of ones, it will represent a simplex in dimension $d-j$, and so its determinant must be some integer $m \leq V(d-j)$. Thus the determinant of N, being the class of σ, is equal to km, and the theorem is proved. □

Corollary 2.2.4. *Since $V(1) = V(2) = 1$ (See Table 1.2), a simplex with an exterior ridge or facet has the same class as that ridge or facet. Thus any simplex with two exterior ridges or two exterior facets must have class 1.*

Lemma 2.2.5. *If σ is a d-simplex ($d \geq 3$)of class c , and if $c > V(d-1)$, then σ has no exterior faces of positive dimension.*

Proof. The formulas in [3] show that for $d \geq 3$ and any $0 \leq j \leq d$, $V(d) \geq V(d-j)V(j)$. Thus if $c > V(d-1)$, then also $c > V(d-2)V(1)$, $c > V(d-3)V(2)$, etc. and by Lemma 2.2.3 no exterior face is possible. □

A cover of the d-cube must, when intersected with a j-face of the cube, induce a cover of that j-cube.

Chapter 3

Previous Work (Summer 2002)

3.1 A Geometric Argument

Smith's bound in dimension 4 is 15, but the minimal triangulation in dimension 4 has size 16. There is conceivably room for a cover smaller than the smallest triangulation. In the Summer of 2002, we proved that this is impossible.

In the 3-cube, all simplices have class 1 or 2 (See Table 1.2). The class-two simplices contain no exterior edges (Lemma 2.2.5), while each class-one simplex contains at most three exterior edges (Lemma 2.2.2). The 3-cube has 12 edges, so by the pigeonhole principle any cover of it must include at least 4 simplices of class 1. If the cover has only 5 simplices total, then the fifth must be class 2 (since the sum of the classes must be at least $3! = 6$).

In the 4-cube, all simplices have class 1, 2, or 3. All simplices with class 2 have at most one exterior facet (Lemma 2.2.4). Also, Cottle ([2]) noted that all the simplices with class 3 share a nonempty mutual intersection. Therefore the cumulative volume of any four simplices is strictly less than $12/4! = 1/2$, and the cumulative volume of any two simplices is strictly less than $6/4! = 1/4$. So it takes at least 5 simplices to cover one half of the 4-volume, and at least 3 to cover a quarter of the 4-volume.

With these facts in mind, we can present two simplified proofs of a result from the Summer:

Theorem 3.1.1 (Bliss, Su 2002). *Mara's triangulation of the 4-cube (i.e., one with eight simplices each of class 1 and class 2) is the only 4-cube cover with fewer than 17 simplices.*

First Proof of Theorem 3.1.1. Let $\mathcal{C}$ be a cover of the 4-cube with fewer than 17 simplices, and let x_i be the number of class-i simplices it contains ($i = 1, 2, 3$). Then,

1. $x_1 \geq 8$. Consider any pair of opposing facets of the 4-cube. The induced cover on each facet must contain at least four class-one 3-simplices. Each of these eight 3-simplices must lie on a different class-1 4-simplex (Lemma 2.2.2).

2. $x_1 < 12$. If $\mathcal{C}$ contained 12 class-one 4-simplices, these would cover at most half the volume of the 4-cube. Since $|\mathcal{C}| < 17$, the remaining half would have to be covered by at most four 4-simplices, which is impossible.

3. $x_2 \geq 4$. Otherwise there would be a pair of opposing facets on which neither induced cover contained a class-two 3-simplex. Then each would have at least six 3-simplices of class one, and each of these would lie on a different 4-simplex of class one, making $x_1 \geq 12$, which was just seen to be impossible.

4. $x_1 < 10$. If $\mathcal{C}$ contained 10 class-one 4-simplices, these would (along with the 4 class-two 4-simplices just mentioned) cover at most three quarters of the 4-cube's volume. The remaining quarter would have to be covered by at most two simplices, which is impossible.

5. $x_2 \geq 8$. Otherwise there would be some induced 3-cube cover with no class-2 3-simplex; it would therefore have at least six class-one 3-simplices. Since the cover on the opposing 3-face must have at least four class-one 3-simplices, and since each of these must lie on a different class-one 4-simplex in $\mathcal{C}$, there would be at least ten class-one 4-simplices in $\mathcal{C}$, which was just proved impossible.

Since $x_1 + x_2 < 17$, $x_1 \geq 8$, and $x_2 \geq 8$, it must be that $x_1 = x_2 = 8$, and the cover is Mara's triangulation. □

Here is another presentation of the same idea:

Second Proof of Theorem 3.1.1. Let $\mathcal{C}$ be any cover of the 4-cube, and let x_i be the number of class-i simplices it contains ($i = 1, 2, 3$). Take any pair of opposite facets of the 4-cube; the induced cover on each must contain at least four 3-simplices of class one. Each of these eight 3-simplices must lie on a different class-one 4-simplex of $\mathcal{C}$, so $x_1 \geq 8$.

Now, consider the eight 3-cube covers induced by $\mathcal{C}$ on the boundary of the 4-cube. One of the following must be true:

Case 1. Each of the 3-cube covers has a class-2 simplex. Then each of these eight class-two 3-simplices must lie on a different class-two 4-simplex in $\mathcal{C}$, and $x_2 \geq 8$. If $\mathcal{C}$ is to have fewer than 17 simplices it must therefore be Mara's.

Case 2. Some pair of covers in opposing cube-facets have no class-two 3-simplex at all. Then each must contain at least six class-one 3-simplices, each of which must lie on a different class-one 4-simplex in $\mathcal{C}$. This is a total of 12 class-one 4-simplices, whose cumulative volume takes up only half of the 4-cube. Since it takes at least five 4-simplices to cover the other half, $\mathcal{C}$ must have at least 17 simplices.

Case 3. Every pair of opposing 3-cube covers contains at least one class-two 3-simplex, and thus $x_2 \geq 4$, but at least one of the 3-cube covers has no class-two 3-simplex. Consider this cover with its opposing one; together they have at least 10 class-one 3-simplices, and so $x_1 \geq 10$. This is a total of at least 14 simplices covering at most three quarters of the 4-cube's volume; it takes at least 3 simplices to cover the remaining quarter so again $\mathcal{C}$ must have at least 17 simplices.

□

3.2 Linear Programming

While geometric arguments like the one just presented produce the sharpest bounds, they tend to exploit special features of the dimension of the cube. They are not, therefore, very good at producing asymptotic results. In an attempt at more generality, I came up with the following idea (which, it later turned out, closely resembles Hughes's work on dissections).

Let $\mathcal{C}$ be a cover of the d-cube, and let x_c be the number of class-i d-simplices in $\mathcal{C}$. Then we can show that a few things have to be true about $\{x_c\}$; in fact these constraints can be formulated as a linear program. If the program is given a (minimizing) objective function of the size of the cover, namely $\sum x_c$, then the optimum value (subject to the constraints) must be a lower bound on the size of any cover of the d-cube: any smaller cover would have to violate one of the constraints!

The most obvious constraint is

$$\sum \frac{i}{d!} x_c \geq 1$$

since the sum of the d-volumes of the simplices must be at least the d-volume of the cube.

In addition, the sum of the $(d-1)$-volumes of all the exterior facets of the d-simplices in $\mathcal{C}$ must be enough to cover each of the facets of the cube. There are $2d$ facets, each with $(d-1)$-volume 1. Each d-simplex has at most one exterior facet, unless its class is 1, in which case it has at most d. In any case the class of the facet will be the class of the simplex, so another bound is

$$\frac{1}{(d-1)!} x_1 d + \sum_{c=2}^{c_{max}} \frac{i}{(d-1)!} x_c \geq 2d.$$

With these two bounds alone, the linear program produced encouraging results. The lower bounds were higher than Smith's lower bounds in dimensions 4, 5, 6, 7, and 9 (See Table 1.1). However, in dimensions 8, 10, and higher, a different set of critical variables took over, and the bounds were worse than Smith's. Clearly, the program needed more constraints to produce good bounds asymptotically.

Chapter 4

A Stronger Program

4.1 The Idea

My first goal for the fall was to improve the bounds given by the program on the previous page. By adding constraints to the program, I could increasing its optimum value, and (I hoped) find an asymptotic benefit over Smith's bounds. I decided to make one constraint for each dimension $d' \leq d$, describing the d'-dimensional "real-estate" that the external d'-faces of the simplices had to cover. There are $2^{d-d'}\binom{d}{d'}$ d'-faces in the d-cube, each with d'-volume 1. So if $F(d, c, d', c')$ describes the maximum number of external class-c' d'-faces on a class-c d-simplex in the d-cube, the following inequalities must hold if the simplices are to cover the cube (and all its faces):

$$\sum_{c=1}^{c_{max}} \sum_{c'=1}^{c'_{max}} \frac{c'}{d'!} F(d, c, d', c') x_c \geq 2^{d-d'} \binom{d}{d'} \quad (d' = 1, 2, \ldots, d) \tag{4.1}$$

However, finding the values of F turned out to be trickier than I anticipated. After much thought, I was able to produce a recurrence relation, which allowed for good computational results in dimensions up to 14, as well as some asymptotic estimates.

4.2 The Recurrence Relation

Let $\Delta(c)$ be the smallest dimension in which a class-c simplex appears. Since $\Delta(c) = \min\{d : V(d) \geq c\}$, calculating values of Δ is just as hard as finding values of V. Some values are tabulated in Table 4.1.

Theorem 4.2.1. *If* $d < \Delta(c)$, $d' < \Delta(c')$, $d' > d$, *or* $c' \not| c$, *then* $F(d, c, d', c') = 0$. *Otherwise,* F *obeys the recursion*

$$F(d, c, d', c') = \sum_{\delta=0}^{d'} \sum_{\gamma=1}^{c'} F(d', c', \delta, \gamma) F(d - d', c/c', d' - \delta, c'/\gamma), \quad (4.2)$$

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

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

Table 4.1: Some values of $\Delta(c)$. Compare Table 1.2

Before we look at a proof outline for the theorem, let's examine some corollaries.

Corollary 4.2.2. *If* (d', c') *is high-class, i.e.* $F(d', c', \delta, \gamma) = 0$ *for* $0 < \delta < d'$, *we have* $F(d, c, d', c') = \min(d/d', \text{weight}(c, c'))$ *(where* $\text{weight}(c, c')$ *denotes the highest power of* c *that divides* c'*).*

Proof. In this case all terms of the summation zero out except $\delta = 0, \gamma = 1$ and $\delta = d', \gamma = c'$, so we have

$$\begin{aligned} F(d, c, d', c') &= \sum_{\delta=0}^{d'} \sum_{\gamma=1}^{c'} F(d', c', \delta, \gamma) F(d - d', c/c', d' - \delta, c'/\gamma) \\ &= F(d', c', 0, 1) F(d - d', c/c', d', c') + F(d', c', d', c') F(d - d', c/c', 0, 1) \\ &= F(d - d', c/c', d', c') + 1 \end{aligned}$$

and by successive application of this recursion, the corollary is shown. □

Corollary 4.2.3. *If* $c' = 1$, *we have* $F(d, c, d', 1) = \binom{d - \Delta(c)}{d'}$.

Proof. This time the only nonzero terms in the sum have $\gamma = 1$, so (proceeding by induction) we have

$$\begin{aligned} F(d,c,d',1) &= \sum_{\delta=0}^{d'} F(d',1,\delta,1)F(d-d',c,d'-\delta,1) \\ &= \sum_{\delta=0}^{d'} \binom{d'}{\delta}\binom{d-d'-\Delta(c)}{d'-\delta} \\ &= \binom{d-\Delta(c)}{d'} \end{aligned}$$

□

Corollary 4.2.4. *If $c' = c$, then $F(d,c,d',c) = \binom{d-\Delta(c)}{d'-\Delta(c)}$.*

Proof. The nonzero terms here are $\gamma = c$, $\delta \geq \Delta(c)$, so (again by induction)

$$\begin{aligned} F(d,c,d',c) &= \sum_{\delta=\Delta(c)}^{d'} F(d',c,\delta,c)F(d-d',1,d'-\delta,1) \\ &= \sum_{\delta=\Delta(c)}^{d'} \binom{d'-\Delta(c)}{\delta-\Delta(c)}\binom{d-d'}{d'-\delta} \\ &= \sum_{d'-\delta=0}^{d'-\Delta(c)} \binom{d'-\Delta(c)}{d'-\delta}\binom{d-d'}{d'-\delta} \\ &= \binom{d-\Delta(c)}{d'-\Delta(c)} \end{aligned}$$

□

Proof sketch for the theorem. If $d > \Delta(c)$, $d' > \Delta(c')$, $d' < d$, and $c'|c$, then we can find a simplex with at least one such exterior face, σ. Project along σ into its orthogonal complement. Then the desired exterior faces fall into three groups: (1) the face we projected across, σ, (2) the faces left in the orthogonal complement, and (3) faces which were destroyed along the way, because they sustained some (positive-dimensional) intersection with groups (1) and (2).

Group (1) has size 1, and group (2) has size at most $F(d-d', c/c', d', c')$. To calculate the size of group (3), we will further subdivide it. Note that the intersection of two exterior faces must be an exterior face. Take a face τ in group (3), and let δ and γ be the dimension and class, respectively, of $\tau \cap \sigma$. The total number of faces τ with such an intersection cannot exceed $F(d', c', \delta, \gamma)F(d-d', c/c', d'-\delta, c'/\gamma)$.

Summing these three groups together, and letting (arbitrarily) $F(d, c, 0, 1) = 1$, we arrive at the above summation formula. □

Using this recurrence and a computer, I generated and solved the linear programs in dimensions 4 through 11; all of the resulting bounds are better than Smith's corresponding ones (See Table 1.1). The optimal solutions are broken down in the appendix.

Chapter 5

Evaluating the Program

5.1 Introduction

The goal of this chapter is to analyze the asymptotic behavior of the bounds produced by the linear program described by (4.1). First, we recast the constraints in a more standard format. Let $M(d)$ be the coefficient matrix defined by

$$M(d)_{d',c} = \sum_{c'=1}^{c} c' F(d, c, d', c'), \quad 1 \leq d' \leq d, 1 \leq c \leq V(d), \tag{5.1}$$

$B(d)$ be the d-element vector defined by

$$B(d)_{d'} = \frac{2^{d-d'} d!}{(d-d')!}, \quad 1 \leq d' \leq d, \tag{5.2}$$

and $N(d)$ be the $V(d)$-element vector defined by

$$N(d)_c = 1, \quad 1 \leq c \leq V(d). \tag{5.3}$$

Consider the linear program whose objective is to minimize $N(d) \cdot \mathbf{x}$ subject to $M(d)\mathbf{x} \geq B(d)$; call the minimum objective $LP(d)$. We have seen that $LP(d)$ is a lower bound on the covering number of the d-cube, and we wish to see how $LP(d)$ grows with d.

Note that if $\mathbf{x}$ is *any* feasible solution to this linear program, then $N(d) \cdot \mathbf{x}$ will give an upper bound on $LP(d)$. Similarly, if $\mathbf{y}$ is a feasible solution to the *dual* program (maximize $B(d) \cdot \mathbf{y}$ subject to $M(d)^T\mathbf{y} \leq N(d)$), then the corresponding objective value will be a *lower* bound on $LP(d)$. We proceed to construct such feasible solutions in the following sections.

5.2 Lower bounds for $LP(d)$

In this section we will exhibit some feasible solutions to the dual linear program, giving lower bounds for $LP(d)$.

To begin, let $\mathbf{y}^0(d) = (1/V(d))\mathbf{e}_d$. Clearly $\mathbf{y}^0(d)$ satisfies the dual program. The objective value corresponding to $\mathbf{y}^0(d)$ is precisely the Euclidean bound shown in Table 1.1. This shows that the bound produced by $LP(d)$ is always as least as good as the Euclidean bound.

Next, let

$$\mathbf{y}^1(d) = \mathbf{y}^0(d) + \left(\frac{1}{v_{d-1}} - \frac{1}{v_d}\right)\mathbf{e}_{d-1}. \tag{5.4}$$

That $y^1(d)$ is a feasible solution to the dual linear program will be established by the following Lemma, which we state without proof:

Lemma 5.2.1. *For all $d \geq 2$ and $k \leq d$,*

$$\frac{1}{V(d-k)} \geq \frac{k}{V(d-1)} + \frac{(1-k)}{V(d)}.$$

Now $\mathbf{y}^1$ will be a feasible solution to the dual program if $M(d)^T\mathbf{y}^1 \leq N(d)$. This means that for any c,

$$\begin{aligned} 1 &\geq \sum_{d'=1}^{d} M(d)_{d',c} y^1_{d'} \\ &= M(d)_{d-1,c} y^1_{d-1} + M(d)_{d,c} y^1_d \\ &= c(d-\Delta(c))\left(\frac{1}{V(d-1)} - \frac{1}{V(d)}\right) + c\frac{1}{V(d)}, \end{aligned}$$

where in the last step we have taken the entries of $\mathbf{y}^1$ from its definition and calculated the relevant entries of $M(d)$ using the Corollaries to Theorem 4.2.1 assuming $c \leq V(d-1)$ (since when $c > V(d-1)$, $M(d)_{d-1,c} = 0$ and the resulting equation has already been verified.) But rearranging the last equation and applying the lemma gives

$$\frac{(d-\Delta(c))}{V(d-1)} + \frac{1-(d-\Delta(c))}{V(d)} \leq \frac{1}{V(\Delta(c))} \leq \frac{1}{c}$$

and thus $\mathbf{y}^1$ is indeed a feasible solution of the dual program. Thus the associated objective value is a lower bound for $LP(d)$; this is simply

$$LP(d) \geq d! \left(\frac{2}{V(d-1)} - \frac{1}{V(d)} \right).$$

Using, again, the bounds from [3] we can estimate this as

$$LP(d) \geq d! 2^d \left(\frac{1}{d^{d/2}} - \frac{1}{(d+1)^{\frac{d+1}{2}\alpha}} \right)$$

but this bound is still worse than Smith's as d becomes very large.

One could construct $\mathbf{y}^2$ in similar fashion, but it would still lose out to Smith in the long run. Constructing $\mathbf{y}^k$ becomes more and more complicated as $k \geq 3$. You would need to get nearly to $\mathbf{y}^d$ before you'd have any chance of beating Smith. In the next section we try to determine if this is even possible.

5.3 An Upper bound on $LP(d)$

Let $\mathbf{u}$ be the vector given by

$$u_{v(d')} = \frac{2^{d-d'} \binom{d}{d'} d'!}{\sum_{c'=1}^{c} c' F(d, v(d'), d', c')}, \quad 1 \leq d' \leq d \tag{5.5}$$

with $u_i = 0$ otherwise. Then $\mathbf{u}$ is a feasible solution to our linear program , and we have

$$\begin{aligned} LP(d) &\leq N(d) \cdot \mathbf{u} \\ &= \sum_{d'=1}^{d} \frac{2^{d-d'} \binom{d}{d'} d'!}{\sum_{c'=1}^{c} c' F(d, v(d'), d', c')} \\ &\leq \sum_{d'=1}^{d} \frac{2^{d-d'} \binom{d}{d'} d'!}{v(d') + \binom{d-d'}{d'}}. \end{aligned}$$

In the last step we have estimated the denominator by throwing away all terms but $c' = 1$ and $c' = V(d')$, and applying Corollaries (4.2.3) and (4.2.2) respectively.

If we use the lower bound for $V(d)$ from [3] and write $\alpha(d) = 1 - \log_d(4/3)$, then we have

$$LP(d) \leq \sum_{d'=1}^{d} \frac{2^{d-d'}\binom{d}{d'}d'!}{(d'+1)^{\binom{d'+1}{2}\alpha(d'+1)} + \binom{d-d'}{d'}}. \tag{5.6}$$

Detailed asymptotic analysis of this inequality should show that our program can never hope to beat Smith asymptotically, no matter how good our numbers are in low dimensions.

Chapter 6

Pebble sets and beyond

6.1 Introduction

In [5], De Loera, Peterson and Su define a *pebble set* of a d-dimensional polytope P as a finite set of points such that each d-simplex of P contains at most one point. The size of a pebble set is clearly a lower bound for the covering number of the polytope. If P has n vertices, then a pebble set of size $(n-d)$ exists; this bound is sharp over all polytopes because for any n and d there is a d-dimensional polytope on n vertices (a so-called "stacked polytope") for which no larger pebble set exists.

However, this $(n-d)$ bound is not always the best possible for a particular polytope. For example, the 3-dimensional octahedron has 6 vertices but can support 4 pebbles (and can be covered by 4 simplices). For the cubes, the bound of (2^d-d) falls far short of even the Euclidean bound (1.1). In this chapter, we will explore the pebble sets of the cubes, introduce a generalization of the pebble set, and examine its relation to simplicial covers.

6.2 Fractional Pebblings and Fractional Coverings

A *(fractional) pebbling* of a d-dimensional polytope P is a signed measure $\mu : P \to \mathbb{R}$ satisfying two criteria: (i) the measure of any subset of a hyperplane spanned by vertices of P is zero, and (ii) the measure of any d-simplex spanned by vertices of P is at most one. If Σ is the set of d-simplices on the vertices of P, we can write these conditions as (i) $\sigma \in \Sigma, B \subseteq \partial\sigma \implies \mu(B) = 0$, and (ii) $\sigma \in \Sigma \implies \mu(\sigma) \leq 1$. The *size* of the pebbling is the measure of the entire polytope, $\mu(P)$.

Two pebblings will be considered equivalent if they assign the same measure to every chamber of P (recall that a chamber is a maximal set disjoint from the boundaries of every d-simplex on P. Thus a pebbling can also be viewed as a real-valued function on the chamber set $X(P)$ of the polytope, or as a finitely supported function on the polytope itself. The pebbling's size is then the sum of the function's values on each of the polytope's chambers.

In complete analogy, we define a *fractional covering* of a polytope P to be a real-valued weight function on the set of simplices Σ of P such that the total weight of all the simplices containing any given point in P is at least one (i.e. $p \in P \implies \sum_{\sigma \in \Sigma, p \in \sigma} w(\sigma) \geq 1$). Note that it suffices to check the total weight of one point in each chamber of P. A *fractional dissecting* is a fractional covering such that every point has a total weight of *exactly* one.

Note that a pebble set as defined in [5] can be seen as a pebbling with range $\{0, 1\}$, and that the size of a fractional pebbling is still a lower bound for the covering number of the polytope. In principle a fractional pebbling may exist with a larger size than any pebble set on the same polytope. Note that $LP(d)$ as defined in Chapter 5 serves as a lower bound for fractional covering size also.

If $\Sigma = \{\sigma_1, \ldots, \sigma_m\}$ are the simplices of P and $X = \{\chi_1, \ldots, \chi_m\}$ are its chambers, we define the m by n *chamber-simplex incidence matrix* $L(P)$ by letting $L_{ij} = 1$ when simplex σ_i contains chamber χ_j, and $L_{ij} = 0$ otherwise.

Consider a linear program with $L(P)$ as the coefficient matrix, and vectors $(1, \ldots, 1)$ as both the objective function and the bound vector. Then the optimal solution of the primal (minimization) program is the size of the smallest fractional covering of P, and the optimal solution of the dual (maximization) problem is the size of the largest fractional pebbling on P. Restricted to an integer program, the primal and dual values give the smallest cover size and largest pebble-set size, respectively.

6.3 Exploiting Symmetry

Now the d-cube is not just any polytope; it has a very special structure with a great deal of symmetry. The following definitions will help exploit that symmetry to greatly simplify the study of pebblings.

Let $S(P)$ denote the group of isometries from P onto itself. (Actually, the functions in $S(P)$ need not be isometries; it is enough that they induce automorphisms on the chamber-simplex hypergraph of P.) A pebbling of P is *symmetric* if it is invariant under composition with any element of $S(P)$. If X is the set of chambers of P, then a symmetric pebbling can be viewed as a function on the quotient set $X' = X/\sim$ (here $\sim$ represents the equivalence relation between a chamber and its image under any element of $S(P)$). A symmetric covering of P is defined the same way; it can be viewed as a function on the quotient set $\Sigma' = \Sigma/\sim$.

Note that, given any pebbling μ of P, there is always a symmetric pebbling μ' with the same size as μ: let $S(P)$ act on the set of all pebblings and set μ' to be the average of the orbit of μ. The same argument applies to coverings.

If we write the sets of isometry families of simplices and of chambers in P as $\Sigma' = \{\sigma'_1, \ldots, \sigma'_{m'}\}$ and $X' = \{\chi'_1, \ldots, \chi'_{n'}\}$, then we can define the m' by n' *symmetrized chamber-simplex incidence matrix* $L'(P)$ by letting L'_{ij} count the number of chambers from family χ'_j present in any simplex of family σ'_i.

The *profile* of a symmetric pebbling is a vector of length n' whose ith element is the measure of any chamber from family χ'_i, and the profile of a symmetric covering is a vector of length m' whose jth element is the weight assigned to any simplex from family σ'_j.

Let Q define the chamber quantity vector; that is, Q_i is the size of the orbit χ'_i. Then if V is the profile of a symmetric pebbling, $Q \cdot V$ will give its size. Further, the condition that no simplex have total measure greater than 1 is exactly the condition $L'V \leq \mathbf{1}$. Thus the size of the largest symmetric fractional pebbling is the objective value to this linear program. The dual program seeks to minimize $\mathbf{1} \cdot W$ subject to $L'^T W \geq Q$. If W is the profile of a symmetric covering, then $\mathbf{1} \cdot W$ gives the covering's size. The bounds in the dual program assure that the total weight assigned to all chambers of a given family must be at least the number of chambers in that family. Since the covering is symmetric, this means that every chamber has total weight at least one. So, just as fractional pebblings are dual to fractional coverings, symmetric fractional pebblings are dual to symmetric fractional coverings.

It is worth noting that the above analysis could have proceeded a different way: we could instead let L'_{ij} count the number of simplices in family σ'_i which contain a given chamber of family χ'_j. Then the quantity vector in question would count the total number of simplices in each orbit, rather than the total number of chambers. The result would be the same: that symmetric fractional pebblings are dual to symmetric fractional coverings.

6.4 Morphs

A fractional pebbling cannot assign values to chambers at random; the chamber weights must relate to each other in a special way. Some of these relations are easy to find, and they can tell us much about the nature of pebblings for particular polytopes. In this section we lay the foundation for such investigations.

Suppose that some subset of the polytope can be dissected in two different ways. Then, if we have a covering of the polytope that includes one of the two dissectings, we may replace it by the other dissecting without changing the total weight assigned to any point in the polytope. Such a procedure we call a *morph*. The difference of the profiles of the two dissectings we call the morph's *profile*.

The easiest example of a morph is a bistellar flip. In this case, the region being dissected is the convex hull of any $(d + 2)$ vertices of the d-dimensional polytope P. We iterate over each of these vertices; each time we leave out the selected vertex and form a simplex from the other $(d+1)$. To determine which dissecting includes the simplex, we first anchor one simplex in one dissecting arbitrarily. After that, every simplex we consider will share a facet with the anchored simplex. If the two vertices not on that facet lie on different sides of it, the simplices go into the same dissecting. If the vertices are on the same side of the shared facets, the simplices go into different dissectings. Degenerate simplices are not counted at all.

If L' is the symmetrized chamber-simplex incidence matrix of P, and if V is the profile of any pebbling and W is a profile of any simplex weighting, then the vector $(L'V)\cdot W$ counts the total pebbling in the region described by W. In particular, if W is the profile of a dissecting, it gives is the size of the pebbling, and if W is a morph profile then the product is 0. We say that one covering is *reachable* from another through a set of morphs if there exists a sequence of the morphs which takes the first covering to the second. Note that a morph need not be applied an integer number of times. In this case the difference of the profiles of the two coverings will be in the span of the profiles of the morphs.

Suppose V is a pebbling profile, and put $U = L'V$. Let M be a matrix whose rows are morph profiles, and let W_0 be any covering profile. Then U must satisfy $MU = \mathbf{0}$ and $I'_m U \leq \mathbf{1}$. Therefore the largest pebbling is no larger than the optimum of the corresponding linear program. The dual of this program seeks to minimize the size of a cover reachable from W_0 through the morphs of M.

There is another way to look at morphs: they are basically dissections of the null set. A morph is a weighting of integers (positive and negative) on the simplices of the polytope such that every point in the polytope gets total weight zero.

6.5 The 3-cube and 4-cube

If P is a d-cube, then $|S(P)| = 2^d d!$, corresponding to the $d!$ permutations of coordinates and the d possible reflections about medial hyperplanes. Thus the symmetrized chamber-simplex incidence matrix is much smaller than the un-symmetrized version! By convention, for the cube we will always number the simplices of Σ so that σ_1 is spanned by

$$\{(0,0,\ldots,0),(1,0,0,\ldots,0),(1,1,0,0,\ldots,0),\ldots,(1,1,\ldots,1,1)\}.$$

We also number the families of Σ' so that $\sigma_1 \in \sigma'_1$. This is convenient because it is possible to dissect the d-cube with $d!$ simplices of this family.

In the d-cube, some chambers cross a medial hyperplane. Such a chamber has orbit size $2^{d-1}d!$; all other chambers have orbit size $2^d d!$ because no isometry can fix them. For simplicity, if a chamber crosses a medial hyperplane we will count it as two chambers, so that all chambers have the same orbit size.

As an example, consider the familiar 3-cube, where the simplices are tetrahedra. In this case the largest pebble set, largest pebbling, smallest covering, and smallest cover all have the same size: 5. In fact this is precisely equal to the $(n-d)$ bound given in [5]. There are 58 non-degenerate simplices in four isometry classes: 24 in σ'_1 (the "loppers", aka "permuters"), 8 in σ'_2 (the "corners"), and 24 in σ'_3 (the "leaners") all have class 1; whereas the 2 in σ'_4 (the "Hadamards") have class 2. There are 144 chambers in three isometry classes: χ'_1 (the "outers"), χ'_2 (the "middles"), and χ'_3 (the "inners"). The symmetrized chamber-simplex incidence matrix is

$$L'(I^3) = \begin{pmatrix} 8 & 8 & 8 \\ 12 & 6 & 0 \\ 4 & 10 & 16 \\ 0 & 24 & 48 \end{pmatrix}$$

The minimal covering has profile $(0,1/2,0,1/2)$; the maximal pebbling has profile $(1/6,0,1/24)$. There are only two bistellar flips that can be performed in the 3-cube; their profiles are $(0,1,-3,1)$ and $(2,-1,-1,0)$.

In the 4-cube there are over 900 chamber families, so we shall not enumerate them. There are only 17 families of simplices; these are listed in table 6.1, along with the 23 bistellar flips that exist. The profiles of these flips span a 12-dimensional space; if V is the profile of a pebbling then $L'V$ must lie in the 5-dimensional orthogonal complement. A linear program can find a pebbling of size 16 in this complement; however it is not clear that this vector of simplex weights is in the column space of L' (i.e. that it corresponds to an actual pebbling).

Family	Quant.	Class	Example
A	16	1	0, 1, 4, 2, 8
B	192	1	0, 4, 2, 8, 9
C	96	1	0, 1, 4, 13, 7
D	384	1	0, 1, 4, 12, 10
E	192	1	0, 12, 14, 8, 15
F	192	1	0, 4, 2, 8, 13
G	384	1	0, 4, 14, 8, 7
H	192	1	0, 1, 2 , 7, 15
I	64	1	0, 4, 2, 8, 15
J	384	1	0, 1, 4, 12, 11
K	192	1	0, 1, 4, 10, 13
L	192	1	0, 1, 12, 10, 15
M	64	2	0, 12, 10, 6, 15
N	64	2	0, 12, 10, 6, 7
O	192	1	0, 8, 5, 11, 7
P	192	2	0, 14, 9, 11, 7
Q	16	3	0, 14, 13, 11, 7

Six Vertices	Relation
0, 1, 2, 3, 4, 8	2B = C+A
0, 1, 2, 3, 4, 12	2E = B+H
0, 1, 2, 3, 4, 13	F+G = J+D
0, 1, 2, 3, 4, 15	2J = K+I
0, 1, 2, 4, 7, 8	3F = M+A
0, 1, 2, 4, 7, 9	B+N = K+2D
0, 1, 2, 4, 7, 11	M+F = 2G+C
0, 1, 2, 4, 7, 15	3H = N+I
0, 1, 2, 4, 8, 15	Q+A = 4I
0, 1, 2, 4, 9, 11	E+F = B+G
0, 1, 2, 4, 9, 14	2J+F = B+L+K
0, 1, 2, 4, 9, 15	P+B = I+2J
0, 1, 2, 4, 11, 15	H+I = F+L
0, 1, 2, 5, 6, 13	H+E = J+D
0, 1, 2, 5, 10, 12	O+2D = 2G+E
0, 1, 2, 5, 10, 15	2J = O+E
0, 1, 2, 5, 11, 14	P+D = L+G+J
0, 1, 2, 5, 11, 15	D+O = H+J
0, 1, 2, 7, 11, 12	Q+2K = C+2P
0, 1, 2, 7, 11, 13	M+2G = C+L+N
0, 1, 2, 7, 12, 13	L+K = G+O
0, 1, 2, 7, 12, 15	P+K = 2O+H
0, 1, 6, 10, 12, 15	Q+N = 3L+M

Table 6.1: Simplices and bistellar flips in the 4-cube. The vertices of the cube are labeled with the integers 0 to 15 by interpreting their coordinate 4-tuple as a binary number.

Chapter 7

Upper Bounds

In this chapter we will establish, by construction, some upper bounds on the covering number (and related numbers) of the d-cube.

7.1 Induced Morphs: Dimensions 1-5

The goal of this section is to establish the following result, first proved in [7].

Theorem 7.1.1. *Let $k_1 = 0$ and $k_d = d(k_{d-1} - 1) + 2^{d-1}$ for $k \geq 2$. Then there is a dissection of the d-cube using $k_d + 2^{d-1}$ simplices.*

In particular, this implies $C(4) \leq 16$ and $C(5) \leq 67$, which are known to be the best possible results in these dimensions. Covers (in fact they are triangulations) of size 16 and 67 for the 4- and 5-cube were described in [6] and [1].

Before proceeding to covers of the d-cube, we first establish two useful formulas:

Lemma 7.1.2.

$$\sum_{i=0}^{d-1} \binom{d-1}{i} \frac{1}{i+1} = \frac{2^d - 1}{d}, and \tag{7.1}$$

$$\sum_{i=0, i \text{ even}}^{d-1} \binom{d-1}{i} \frac{1}{i+1} = \frac{2^{d-1}}{d}. \tag{7.2}$$

Proof. These formulas can be obtained by expanding $(1+x)^{d-1}$ and $((1+x)^{d-1}+(1-x)^{d-1})/2$ using the binomial formula, integrating once, and evaluating at $x = 1$. □

Equation (7.2) shows that the even terms contribute slightly more than the odd terms in the sum of equation (7.1).

Our construction of a small cover for the d-cube will make use of a sequence of bistellar flips. To explain these, we first define the simplex family ${}^c\sigma_i^d$ for integers d and i (the c is just a label), as follows: let p_0 be a point of taxicab distance i from 0 in a $(d-1)$-cube (as a facet of the d-cube), and let $p_1, \ldots, p_{d-1}$ be the neighbors of p_0 in the same $(d-1)$-cube. Let q be the neighbor of 0 that is *not* in that $(d-1)$-cube. Then ${}^c\sigma_i^d$ is the isometry family of the d-simplex defined by the points $\{p_0, p_1, \ldots, p_{d-1}, q\}$. Note that ${}^c\sigma_0^d$ is the isometry class of the corner in dimension d.

Lemma 7.1.3 (The $(i+1)$-fold corner-flip.). *The following relation defines a morph for every d and i:*

$$(i+1)^c\sigma_i^d = {}^c\sigma_0^d + {}^u\sigma_i^d$$

where ${}^u\sigma_i^d$ is some d-simplex and ${}^u\sigma_0^d$ is the empty set.

Proof. For $i = 0$ the equation is trivial; therefore fix $i \geq 1$. Let the vertices of ${}^c\sigma_i^d$ be labeled as above, and order the p_j in such a way that $\{p_1, \ldots, p_i\}$ are the neighbors of p_0 which lie in the same i-cube as 0. Let r be the neighbor of p_0 which does not lie in the same $(d-1)$-cube as 0. The morph comes from a bistellar flip on the vertices $\{p_0, p_1, \ldots, p_{d-1}, q, r\}$. The exclusion of p_0 forms the simplex we'll call ${}^u\sigma_i^d$. The exclusion of any of $p_1, \ldots, p_i$ forms a copy of ${}^c\sigma_i^d$, as does the exclusion of r. The exclusion of any of $p_{i+1}, \ldots, p_{d-1}$ makes a degenerate simplex, since the $(i+3)$ points $\{p_0, \ldots, p_i, q, r\}$ all lie in the same $(i+1)$-cube. Finally, the exclusion of q gives the corner simplex on p_0, which is the desired ${}^c\sigma_0^d$. $\square$

These morphs allow us to recursively build a cover for the d-cube, as the next theorem shows.

Theorem 7.1.4. *Suppose the $(d-1)$-cube has a covering (dissecting) C which assigns weight $1/2$ to every corner, and total weight x to all other simplices. Then the d-cube has* *a covering (dissecting) C' which assigns weight $(1/2 - 2^{-(d+1)})$ to every corner and total weight $d(x-1/2) + 2^{d-1} - 1/2$ to all other simplices.*

Proof. We begin by dissecting the d-cube with $d!$ simplices (the "permuters"), each of which includes 0 and its antipode. This induces a cover on each of the d facets which include 0; each of these covers uses $(d-1)!$ simplices. By hypothesis we can morph these $(d-1)!$ simplices into the simplices of C. Each $(d-1)$-dimensional corner-simplex in the $(d-1)$ cube which is taxicab distance i from 0 will be coned over, producing a copy of ${}^{c}\sigma_i^d$. There are $\binom{d-1}{i}$ such corners, each with weight $1/2$. This applies equally to each of the d facets, so the covering of the d-cube will have total size

$$dx + \frac{d}{2}\sum_{i=0}^{d-1}\binom{d-1}{i},$$

where each term in the summation corresponds to a simplex in family ${}^{c}\sigma_i^d$. But applying the $(i+1)$-fold corner-flip and formula (7.1), we see that

$$\begin{aligned}\sum_{i=0}^{d-1}\binom{d-1}{i}{}^{c}\sigma_i^d &= \sum_{i=0}^{d-1}\binom{d-1}{i}\frac{1}{i+1}\left({}^{c}\sigma_0^d + {}^{u}\sigma_i^d\right)\\ &= \frac{2^d-1}{d}{}^{c}\sigma_0^d + \left(\frac{2^d-1}{d} - 1\right)\text{ other simplices}\end{aligned}$$

where in the last line we have grouped together all the ${}^{u}\sigma_i^d$, remembering that ${}^{u}\sigma_0^d$ is empty. Since there are a total of 2^d corners (i.e. ${}^{c}\sigma_0^d$) and this cover uses $2^{d-1} - 1/2$ of them, after symmetrizing the weight assigned to any one of them will be $1/2 - 2^{-(d+1)}$, and the total weight of the remaining simplices is $dx + \frac{d}{2}\left(\frac{2^d-1}{d} - 1\right) = d(x-1/2)+2^{d-1}-1/2$, as desired. Clearly if the C was a dissection then the resulting cover is also. □

This is a nice construction, but unfortunately it can't be iterated because the weight on the corner-simplices is slightly lowered. It is actually possible to collect all the ${}^{u}\sigma_i^d$ simplices and combine them in complicated ways to produce the remaining $1/2$ corner simplex, and in the process reduce the total size of the cover.

However, there is an easier way! Recall equations (7.1) and (7.2). They show that the $(d-1)$-dimensional corners on points which are an even distance from the origin (called the "even corners") morph a little more strongly than the other ones (the "odd corners"). This is because we chose the $d!$ permuters to include 0 and its antipode in the d-cube. If we could require the original covering to have weight 1 on the even corners and 0 on the odds (rather than $1/2$ on all of them), we'd get a slightly stronger result—but a symmetric covering cannot distinguish between even and odd corners (which are isometric).

Therefore, if $S(d)$ is the group of the $2^d d!$ isometries of the d-cube, consisting of all permutations on the coordinates and reflections about medial hyperplanes, let $S'(d)$ be the subgroup of the $2^{d-1}d!$ isometries which use *only an even number of reflections*. A pebbling or covering of the d-cube will be called *semisymmetric* if it is invariant under composition with the elements of $S'(d)$. When S' acts on the set of d-simplices, the corner-simplices will break down into two orbits, called the *even corners* (which include the corner on 0) and the *odd corners*. A semisymmetric covering or dissecting which assigns weight 1 to the even corners and 0 to the odd corners will be called *corner-cutting* as in [8]. Note that the $(i+1)$-fold corner-flip produces an even corner exactly when i is even, and none of the ${}^u\sigma_i^d$ could be an odd corner.

Now we can restate Theorem 7.1.4 with a stronger hypothesis and a stronger conclusion:

Theorem 7.1.5. *Suppose the $(d-1)$ cube has a semisymmetric corner-cutting covering (dissecting) C with total weight $2^{d-2}+x$. Then the d-cube has a semisymmetric corner-cutting covering (dissecting) with total weight $2^{d-1}+d(x-1)+2^{d-1}$.*

Proof. We follow exactly the same procedure as in the proof of theorem 7.1.4. This time, rather than putting weight $1/2$ on each ${}^c\sigma_i^d$, C will put weight 1 on those with even i and 0 on those with odd i. Thus the induced cover will have the form

$$dx \text{ simplices} + d \sum_{i=0, i \text{ even}}^{d-1} \binom{d-1}{i} {}^{c}\sigma_i^d$$

and after applying the $(i+1)$-fold cover-flip we will have

$$\begin{aligned}
\sum_{i=0, i \text{ even}}^{d-1} \binom{d-1}{i} {}^{c}\sigma_i^d &= \sum_{i=0, i \text{ even}}^{d-1} \binom{d-1}{i} \frac{1}{i+1} \left({}^{c}\sigma_0^d + {}^{u}\sigma_i^d\right) \\
&= \frac{2^{d-1}}{d} \text{ even corners} + \left(\frac{2^{d-1}}{d} - 1\right) \text{ other simplices.}
\end{aligned}$$

So after semisymmetrizing, we will have 1 on all the even corners, 0 on all the odd corners, and a total weight of $dx + 2^{d-1} - d$ on the other simplices, as desired. Again, it is clear that if C was a dissecting, the result is also. □

Happily this result can be applied inductively, starting with the usual triangulation of the 2-cube (which happens to be a semisymmetric corner-cutting dissection), giving Theorem 7.1.1. The x_d defined there can be given in closed form:

$$x_d = 1 + d! \sum_{i=4}^{d} \frac{2^{i-1} - 1}{i!}$$

and as $d \to \infty$ the summation term converges to $(e^2 - 2e - 1)/2 \approx 0.47625$, so we have the asymptotic estimate that

$$C(d) \leq 2^{d-1} + 1 + d!(0.47625). \tag{7.3}$$

7.2 *Morph Products: Dimensions* 6 *and up*

For all its strength in low dimensions, the construction in the previous section gives a very weak asymptotic estimate. The purpose of this section is to find a better asymptotic upper bound for the covering number of the d-cube.

Let σ_n^d and τ_n^d be the isometry families of the d-simplices represented by the matrices

$$\begin{pmatrix} I_n & 0 \\ 1 & L_{d-n} \\ 0\ldots0 & 0\ldots0 \end{pmatrix}, \quad \begin{pmatrix} I_n & 0 \\ 1 & L_{d-n} \\ 1\ldots1 & 0\ldots0 \end{pmatrix}, \tag{7.4}$$

respectively (L_n denotes the n by n lower-triangular matrix of 1s). As special cases, observe σ_1^d is the "permuter" simplex, σ_d^d is the "corner" simplex, and τ_d^d is the "leaner" simplex (consisting of a vertex and all the neighbors of its antipode).

By adding the vector $\mathbf{e}_n = (0,\ldots,0,1,0,\ldots,0)$ to σ_{n-1}^d and performing a bistellar flip, we find the relation

$$n\sigma_{n-1}^d = \sigma_n + \tau_n. \tag{7.5}$$

Applying this recursively yields

$$\sigma_1^d = \frac{1}{d!}\sigma_d^d + \sum_{n=2}^{d} \frac{1}{n!}\tau_n^d. \tag{7.6}$$

Starting from the standard cover of $d!\sigma_1^d$, we can thus achieve

$$C(d) \leq 1 + d!\sum_{n=2}^{d}\frac{1}{n!} \leq 1 + (e-2)d! \tag{7.7}$$

but this is not good enough asymptotically.

However, if we notice that

$$\tau_n^d = \tau_n^n \times \sigma^{d-n} \tag{7.8}$$

as an orthogonal product, then we can recycle the τ simplices to make an even smaller cover. It is not clear that morphs which are valid on the faces of the simplices will still be valid on the simplices themselves, but if this were true we could proceed in this way:

$$\sigma_1^d = \frac{1}{d!}\sigma_d^d + \sum_{n=2}^{d}\frac{1}{n!}\left(\tau_d^d \times \sigma_1^{d-n}\right) \tag{7.9}$$

In this way we can reduce every σ_1^k to a sum of simplices involving only products of τ_i^i and σ_j^j. If we let s_k denote the total number of simplices that result from this process, and put $c_k = k!s_k$, we will have the relations

$$s_d = \frac{1}{d!} + \sum_{n=2}^{d} \frac{1}{n!} s_{d-n} \tag{7.10}$$

$$c_d = 1 + \sum_{n=0}^{d-2} \binom{d}{n} c_n. \tag{7.11}$$

Taking $s_2 = 1$ since the 2-dimensional permuter is already a corner, we get the sequence of cover sizes (starting with $d = 0$)

$$1, 1, 2, 5, 18, 77, 408, 2473, 17342, \ldots \tag{7.12}$$

which is not very impressive in low dimensions. However, though I have not yet found a closed form for the c_d, computation shows that they grow as

$$c_n \approx (0.8724532496)^n n! \tag{7.13}$$

which is only slightly worse than the $0.840463^n n!$ which Smith cites as the best known asymptotic upper bound for triangulations, using Haiman's recursive result and the computed value $T(7) \leq 1493$. The latter bound is thus only valid when n is a multiple of 7, whereas the previous bound holds for all n.

Appendix A

Constructing and Solving the Program in Low Dimensions

Here is the `perl` code used to generate the linear program. The output format is suitable for piping to `lp_solve`, a linear program solver.

```
#!/usr/bin/perl
$d = shift;
print "min: ";
for ($c=1; $c <= maxvol($d); $c++) {
    print "+y$c ";
}
print ";\n";
for ($dp=1; $dp<=$d; $dp++) {
#for ($dp=$d-2; $dp<=$d; $dp++) {
    print "dim$dp: ";
    for ($c=1; $c <= maxvol($d); $c++) {
        $coef = 0;
        for ($cp =1; $cp <= maxvol($dp); $cp++) {
            $coef += $cp * bound($d,$c,$dp,$cp);
        }
        print "+$coef y$c ";
    }
    print " >= ". (fact($dp)*2**($d-$dp)*choose($d,$dp)). ";\n";
}
```

```
sub fact {
    my $k = shift;
    my $returnval = 1;
    while ($k > 0) {
        $returnval *= $k--;
    }
    return $returnval;
}

sub choose {
    my $n = shift;
    my $k = shift;
    my $returnval = 1;
    while ($k > 0) {
        $returnval *= $n-- / $k--;
    }
    return $returnval;
}

sub maxvol {
    my @vals = (1,1,1,2,3,5,9,32,56,144,320,1458,3645,9477);
    return $vals[shift];
}

sub mindim {
    my $dim = shift;
    my $returnval = -1;
    while (maxvol(++$returnval) < $dim) {}
```

```
    return $returnval;
}
sub bound {
    my ($d,$c,$dp,$cp) = @_;
    if ($dp > $d) {
        return 0;
    }
    my $div = $c/$cp;
    if ((int($div) != $div) || $div > maxvol($d-$dp)) {
        return 0;
    }
    if ($dp == $d) {
        if ($cp ==$c) {
            return 1;
        }
        return 0;
    }
    if (($c > maxvol($d)) || ($cp > maxvol($dp))) {
        return 0;
    }
    if ($dp == 0 && $cp == 1) {
        # Special case, for easy computation
        return 1;
    }

    my $returnval=0;
    for (my $delta = 0; $delta <= $dp; $delta++) {
        for (my $k=1; $k <= $cp; $k++) {
```

```
                $returnval += bound($dp, $cp, $delta, $k)*
                    bound($d-$dp, $c/$cp, $dp-$delta, $cp/$k);
            }
        }
        return $returnval;
    }
```

When the output of the above code is piped to `lp_solve`, a linear program solver, the results come out as in Table A.1. There are several interesting patterns to be noted here. Bear in mind that the numbers do not represent *actual* covers; they are merely an idealized set of simplices which minimize the objective function relative to the linear constraints. Also, there is no guarentee of uniqueness; for example an alternate minimum exists in dimension 4 at $x_1 = 8$, $x_2 = 8$, $x_3 = 0$ (which actually corresponds to a possible cover, unlike the $(12, 4, 0)$ listed in the table).

- Class-one simplices are always useful. They can have many class-one exterior faces of low dimension, and these do not multiply the top-dimension simplex's class to high numbers.

- While the highest available class increases very rapidly as the dimension increases, these high-class simplices are never used, because they can't have any exterior faces at all. The highest simplex that is actually *useful* seems to grow at a much more modest rate

- Nonetheless, the only simplices ever used in these minimal values are simplices whose class are maximal for *some* dimension. Compare Table 1.2. Yet, the class-nine simplex is never used.

- The class-144 simplex seems to do an abnormally good job of compromising a high class with a very composite one (so as to provide many exterior low-dimensional faces).

Class\Dim	4	5	6	7	8	9	10	11	12	13 [1]
Total	16	60	249.6	116.9	4680	21384	88172	494547	2681790	14921900
1	12	20		288	2880	10368	12406.2	32256	246917	118491
2	4	20	120							
3		20								
5			57.6	806.4						
9										
32				22.5	1800	11016		332640	591113	2582520
144							3015.38	98560	1742020	10082300
320							2965.85	26611.2		
1458								4480	74013.4	1710600
9477										109447

Table A.1: Simplices used in optimal LP solutions

Bibliography

[1] M.N. Broadie and R.W. Cottle. A note on triangualting the 5-cube. *Discrete Math*, 52:39–49, 1985.

[2] Richard W. Cottle. Minimal triangulation of the 4-cube. *Discrete Math.*, 40(1):25–29, 1982.

[3] Matthew Hudelson, Victor Klee, and David Larman. Largest j-simplices in d-cubes: Some relatives of the hadamard maximum determinant problem. *??*, 1996.

[4] Robert B. Hughes. Lower bounds on cube simplexity. *Discrete Mathematics*, 133:123–138, 1994.

[5] Jesus A De Loera, Elisha Peterson, and Francis Edward Su. A polytopal generalization of sperner's lemma. *J. Combinatorial Theory Ser. A*, 100(1):1–26, October 2002.

[6] Patrick Scott Mara. Triangulations for the cube. *J. Combinatorial Theory Ser. A*, 20(2):170–177, 1976.

[7] John F. Sallee. A triangulation of the n-cube. *Discrete Math*, 40(1):81–86, 1982.

[8] Warren D. Smith. A lower bound for the simplexity of the n-cube via hyperbolic volumes. *Europ. J. Combinatorics*, 21:131–137, 2000.

[9] M. J. Todd. The computation of fixed points and applications. In *Lecture Notes in Economics and Mathematical Systems*, New York, 1976. Springer-Verlag.