Counterexamples in Discrete Geometry

by

Huntington Tracy Hall

B.S. (Brigham Young University) 1996

M.S. (Brigham Young University) 1998

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Mathematics

in the

GRADUATE DIVISION
of the
UNIVERSITY of CALIFORNIA, BERKELEY

Committee in charge:

Professor Robion Kirby, Chair
Professor Vaughan F. R. Jones
Professor Jonathan Shewchuk

Fall 2004

Counterexamples in Discrete Geometry

Abstract

Counterexamples in Discrete Geometry

by

Huntington Tracy Hall
Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Robion Kirby, Chair

This dissertation provides answers to two separate questions in the field of discrete geometry. The two classes of objects under consideration are constructed according to simple rules, but in both cases the objects are shown to have a richer structure than had been thought possible.

Firstly, Anders Björner and Kimmo Erikkson, generalizing a conjecture of Robert Samuel Simon, asked in 1991 whether the independence complex of a matroid is always extendably shellable. The simplest non-trivial test case for the question is realized geometrically as the dual of a cube of any dimension. It is here shown that the cross polytope, contrary to expectation, is not extendably shellable in dimension 12 or higher. A key part of the argument includes the construction of a uniform oriented matroid of rank 4 with a mutation-free element. Although oriented matroids with this property have been constructed before, the example here presented is smaller (13 elements as opposed to 17 or 21) than was previously known to exist.

Secondly, investigations into the complexity of certain calculations involving Gröbner bases had led Bernd Sturmfels and Rekha R. Thomas in 1997 to a conjecture which was subsequently reduced to a question of plane geometry: If a convex lattice polygon is subdivided along all possible lattice line segments, is there a uniform bound on the number of edges in a chamber of the resulting complex? A stronger form of the conjecture quickly failed with the discovery of a lattice polygon whose chamber complex contained a pentagon, and a negative answer to the lattice polygon question would disprove the conjecture entirely. Computerized searches have since found chambers with as many as 15 edges, but without suggesting any pattern of lattice polygon chambers increasing beyond bound. It is here shown that chambers exist, not only of all sizes, but of all shapes: Any convex region with point symmetry is affinely approximated, to arbitrary precision, by chambers of some sequence of lattice polygons.

Professor Robion Kirby
Dissertation Committee Chair

Contents

Chapter 1

Introduction

The results of this dissertation answer open questions about two rather different sorts of objects: the cross polytopes, of which there is exactly one example in each dimension; and lattice polygons, of which there are infinitely many examples, all of them having dimension 2. Both classes of objects come from the realm of discrete geometry. That is to say, they are geometric objects, but objects which in every case can be fully described by a finite (and often fairly small) collection of combinatorial data.

Objects described by such a limited collection of data have a correspondingly rigid structure, and investigation of small examples often suggests limits on the sort of complexity that any such example might be able to exhibit. In Chapter 2 we investigate the cross polytope, which had been conjectured to be extendably shellable in every dimension, or in other words never to have sufficent complexity to get tangled up in itself during the process of shelling. In Chapter 3 we investigate the set of lattice polygons, which had been conjectured only to have a bounded degree of complexity, as measured by the number of edges of the tiny chambers that are obtained when a lattice polygon is split simultaneously along all of its lattice line segments. In both cases, we discover that a richer structure exists than had been supposed possible.

In dealing with objects described by combinatorial information, it is often useful to find some way of representing the same information in a completely different setting, or in other words to make use of the discrete nature of the problem rather than using the geometry. This is the case for our first result; we find a surprising connection between shellings of the cross polytope and pseudoplane arrangements and related oriented matroids. The final obstruction to extendable shellability is a basic phenomenon borrowed from the field of topology: In dimensions 3 and greater, there is such a thing as a knot.

The second problem, by contrast, is solved by direct investigation of the geometry, and indeed by pretending that the description of the object is not discrete at all but rather continuous. The resolution of this problem also hinges on a basic fact coming from a seemingly unrelated area of mathematics, in this case number theory: When all the denominators are low, rational numbers are poor approximations of other rational numbers.

Chapter 2

The cross polytope is not extendably shellable

2.1 Background

A cell complex is said to be of "pure dimension" d if it is equal to the union of its d-dimensional cells, with no leftover lower-dimensional pieces. Such a cell complex is called "shellable" if it can be built up inductively by attaching d-dimensional cells along topological balls or spheres of dimension $d-1$. Several important classes of cell complexes can be shown to have this property. For example, the boundary of a convex polytope of dimension $d+1$ is a cell complex of pure dimension d, and it was long thought that such complexes must be shellable (Schläfli assumes it implicitly in his 1852 calculation [Sch50] of Euler characteristic), a fact which was finally proven by Bruggesser and Mani [BM71], and which is indispensible to much of what is now known about convex polytopes, such as McMullen's proof [McM70] of the Upper Bound Theorem.

Another important type of easily shellable complex is given by the independence complex of a matroid. A matroid is simply a finite ground set $\mathcal{M}$ and some reasonable definition of independence, where by "reasonable" we mean that a subset of an independent set is independent, and that for any subset A of $\mathcal{M}$, all the maximal independent sets are of the same size, rank(A). (To be pedantic, we also require that at least one set, namely the empty set, be independent—which is only reasonable.) The most common examples of $\mathcal{M}$ are

1. The columns of a matrix, where (linear) independence and rank preserve their usual meanings, or

2. The set of edges of a graph, for which independence is defined to mean an absence of cycles, and for which the rank of a subset is the number of vertices minus the number of connected components.

Given a matroid $\mathcal{M}$, the independence complex of $\mathcal{M}$ is the abstract simplicial complex whose simplices correspond to the independent sets of $\mathcal{M}$. More concretely, the independence complex of $\mathcal{M}$ may be constructed by associating the n elements of $\mathcal{M}$ with n affinely independent vectors, such as the n standard basis vectors of $\mathbb{R}^n$. The convex hull of a subset of these points is a simplex of the independence complex of $\mathcal{M}$ if and only if the corresponding set of elements of $\mathcal{M}$ is independent. It was shown in the dissertation of J. Scott Provan, and also reported by Provan and Billera in [PB80], that the independence complex of any matroid is shellable. This implies, for example, that the independence complex of any matroid has the homotopy type of a wedge of spheres.

At any stage of a shelling, the criterion for attaching the next cell is easy to verify. (This is particularly true in the case of a simplicial complex: It is required only that each d-simplex in the sequence be attached along a union of its facets.) Any sequence of d-cells in the complex satisfying this criterion is called a "partial shelling", whether or not the sequence includes every cell of the complex, and regardless of whether it is part of a complete shelling.

Given a partial shelling which we wish to extend by a single cell, it is not difficult either to find a new cell satisfying the criterion or to show that none exists given the choices that have already been made. This does not mean, however, that it is always a simple matter to determine whether a particular partial shelling can be completed to an entire shelling, let alone to determine whether the entire complex itself is shellable starting from nothing. It may be possible to reach an impasse because of choices made at various earlier stages of the partial shelling, even though each of those choices satisfied the criterion for attachment at the time the choice was made. What is more, it may be impossible, short of an exhaustive search, to determine in what way, if any, the earlier choices can be corrected.

In many cases this difficulty—the global consequences of local choices—is not an issue, because all choices are made from a global perspective. In the shellability proofs already mentioned, for example, the complex under consideration has a global structure, coming from a polytope or from a matroid, which is exploited to specify an explicit ordering on all the cells simultaneously such that the attachment criterion is always satisfied. There are applications, however, where there is not enough structure to divine a global ordering *a priori*, leaving no choice but to build the shelling in stages. Given the large number of viable choices at each stage, it is infeasible to follow through with all possibilities, and it would be useful in such cases to have a guarantee that the choices do not matter—that any cell satisfying the criterion for attachment can be added without interfering with the remainder of the process. A complex for which there is such a guarantee is called "extendably shellable" [DK78b] because every partial shelling can be extended to a shelling of the entire complex.

Which cell complexes, then, are extendably shellable? It might be hoped that certain cell complexes which satisfy a large number of other "nice" properties would satisfy this property as well. Let us consider the two classes of complexes already mentioned, both of which are at least known to be shellable: boundaries of convex polytopes and independence complexes of matroids.

The process of shelling the boundary of a polytope is something that can be pictured without too much difficulty, at least in low dimensions. A 3-polytope, or polyhedron, has a boundary homeomorphic to an ordinary sphere, and the process can be visualized as one of claiming territory one patch at a time across an entire globe: each annexed region must share a single border with the territory already occupied. The boundary of a 4-polytope is homeomorphic to the 3-sphere, which can be visualized as ordinary space together with a "point at infinity"—a universe in which every direction eventually brings you back to where you started—and the 3-cells of the boundary complex can be pictured by dividing the universe into finitely many regions with a coarse foam of soap bubbles. The game, once again, is to conquer the entire universe by laying claim to one region at a time, such that each added region abuts the occupied territory along a single simply-connected surface.

Bruggesser and Mani [BM71] accomplish this process, in any dimension, by sweeping "evenly" across the boundary of the polytope in the following way: Pick a point in the interior of the polytope, and a generic line intersecting that point. This defines a pair of opposite rays; move outward along one of them. Once outside the polytope, record the order in which the facets on one half of the boundary become "visible". Then return along the opposite ray, recording the order in which the remaining facets disappear from view. The sequence of facets so obtained is a shelling order for the boundary of the polytope.

Is it really necessary, though, to be so careful about sweeping across the boundary "evenly"? In dimension 2—that is, for the boundary of a polyhedron—the answer is: No, we can be careless. Any scheme whatsoever will succeed, so long as each added region is contiguous with the occupied territory along a single frontier. The argument goes along the following heuristic lines: Of those regions which border along at least one frontier, some may border along additional frontiers or at additional points. Removing these problematic regions breaks the unoccupied territory into smaller subdomains, not all of which are empty, and any minimal remaining piece is a cul-de-sac with at least one region that can legally be adjoined. A more rigorous proof is given by Danaraj and Klee [DK78a].

In dimension 3, however—that is, for the boundary of a 4-polytope—the situation is more subtle. The crucial difference appears to be the phenomenon of knotting: The boundary of a 4-polytope can fail to be extendably shellable if there are enough cells (if the subdivision is fine enough, that is) for a partial shelling to follow the path of a knotted graph. The obstructions encountered in this way—that is, the set of 3-cells remaining when the impasse is reached—are thickened versions of non-collapsible 2-complexes. (A well-known example of such a 2-complex is Bing's house with two rooms.) The hope that every polytope boundary might be extendably shellable thus proves to be overly optimistic, and indeed in some sense "most" polytope boundaries are not extendably shellable [Zie98, HZ00]. The question remains relevant for certain classes of "small" polytopes, however, which appear to be coarsely divided enough to prevent knots from appearing. One test case of particular interest is the boundary of the cross polytope, or generalized octahedron, a simplicial polytope which is the dual, in any dimension, of the n-cube. The cross polytope can be realized as the convex hull of $2n$ vectors in $\mathbb{R}^n$: the n standard basis vectors and their negatives.

The second class of candidates for the property of extendable shellability is that of independence complexes of matroids. Provan and Billera, in showing [PB80] that such complexes are always shellable, actually show something stronger: They describe a graded series of properties, the strongest of which is "vertex decomposability" and the weakest of which is equivalent to shellability, and they show that the independence complex of a matroid is vertex decomposable. There was some question as to whether this stronger property might on its own imply extendable shellability, which was recently answered [MT03] in the negative—indeed, neither property implies the other. It is true, however, that simplices of a matroid complex can be added to the shelling "greedily" in lexicographic order, without regard to what remains to be shelled. Simon has conjectured [Sim94] that the independence complex of a uniform matroid (or in other words, the k-skeleton of an n-simplex) is always extendably shellable, and Björner and Eriksson have extended [BE94] this question to matroids in general. They also answer the question in the affirmative whenever the complex is of dimension 2, or in other words for all matroids of rank 3. This is analogous to the case of polytope boundaries, where the 2-dimensional boundary complex of a 3-dimensional polyhedron is always extendably shellable. Unlike the case of polytope boundaries, however, the general question for matroid complexes remains open; there is no known example of a matroid whose independence complex is not extendably shellable. Perhaps the simplest infinite family of matroids to which this question could be applied are the matroids consisting of n pairs of parallel elements, represented by a matrix of the form $[I|I]$ or by the graph of a tree with every edge doubled. An obvious test case for the question, then, is the independence complex of this matroid, which has a nice geometric realization—so nice, in fact, that we have already mentioned it: the boundary of the cross polytope.

The main result of this chapter is that even in such a simple family of examples as the cross polytopes of all dimensions, any hope of extendable shellability is once again overly optimistic: The cross polytope is not extendably shellable in dimension 12 or higher. (We will speak interchangeably of shellings of the cross polytope or of its boundary complex, because any shelling of one is combinatorially equivalent to a shelling of the other, assuming the cross polytope is triangulated with a single interior vertex at the center.)

The proof exploits a correspondance between partial shellings of the cross polytope and generic affine pseudoplane arrangements, which are closely related to uniform oriented matroids. In particular, the obstruction is obtained by constructing a rank-4 uniform oriented matroid with a mutation-free element, which is a smaller example (13 elements instead of 17 or 21) than was previously known.

2.2 Preliminaries

2.2.1 Shellings and shellability

Certain naturally arising cell complexes can be shown inductively to have the homotopy type of a wedge of spheres.

Let $\mathcal{C}$ be a finite cell complex such that each maximal cell is of dimension d; such a cell complex is said to be *pure* of dimension d. A *shelling* of $\mathcal{C}$ is a total ordering

$c_1, c_2, \ldots$ of the maximal cells such that, for each $k > 1$,

$$c_k \cap \bigcup_{j=1}^{k-1} c_j$$

is either a ball of dimension $d-1$ or a sphere of dimension $d-1$ (in which case the intersection is the entire boundary of c_k). In the first case, c_k retracts onto the previous subcomplex; in the second case, attaching the new cell is homotopically equivalent to attaching a d-sphere along a single point. The entire complex is thus homotopic to a finite wedge of d-spheres.

A subcomplex of the boundary of a d-simplex is homeomorphic to either a ball or a sphere of dimension $d-1$ if and only if it is pure of dimension $d-1$. This simplifies the criterion in the case that $\mathcal{C}$ is a pure simplicial complex of dimension d: A shelling of C is then a total ordering $c_1, c_2, \ldots$ of the maximal simplices such that, for each $k > 1$,

$$c_k \cap \bigcup_{j=1}^{k-1} c_j$$

is a non-empty union of facets of c_k.

2.2.2 Extendable shellability

A much stronger notion is that of an *extendably shellable* complex, in which no sequence of legal moves can lead to an impasse, and every partial shelling can be extended to a complete shelling. To be precise, a complex C of pure dimension d is extendably shellable if every shellable proper subcomplex D of C intersects some cell of $\overline{C \setminus D}$ along a ball or sphere of dimension $d-1$.

Björner and Eriksson [BE94] ask whether the independence complex of any (non-uniform) matroid is always extendably shellable. Perhaps the simplest test case is the rank-n matroid consisting of n pairs of parallel elements, representing the linear dependencies of the $2n$ columns of the matrix $[I|I]$, or, as a graph, representing the cycle dependencies of a tree in which every edge has been doubled. The independence complex of this matroid is the boundary of the *cross polytope,* or generalized octahedron. The cross polytope is the dual to the n-dimensional cube, and can be constructed as the convex hull of the n unit vectors in $\mathbb{R}^n$ together with their negatives. The sequence of boundaries of generalized octahedra provides arguably the simplest infinite sequence of simplicial complexes of increasing dimension (besides the trivial example of generalized tetrahedra, i.e. a simplex itself or its boundary); each object in the sequence is the suspension of the previous one, starting with a pair of points. It is natural to ask whether such an uncomplicated object is, in each dimension, extendably shellable; the main result of this chapter will answer that question in the negative.

2.2.3 Shellings of the cross polytope

Let $(e_1, \ldots, e_n)$ represent the coordinate unit vectors in $\mathbb{R}^n$. The cross polytope of dimension n is the convex hull of the $2n$ points $\{e_1, -e_1, \ldots, e_n, -e_n\}$, and each of the 2^n facets corresponds to a binary choice, for each $1 \leq i \leq n$, between a standard basis vector $(-1)^0 e_i$ and its negative $(-1)^1 e_i$. We have written these vectors in such a way as to suggest the correspondence to binary strings of length n; the facet indicated by a list of n exponents from the set $\{0, 1\}$ is the convex hull of the corresponding vectors. An ordering of the boundary complex of the cross polytope thus corresponds to a list $(S_1, S_2, \ldots, S_{2^n})$ of the 2^n binary strings of length n, and we can describe in purely combinatorial terms what it means for such an ordering to constitute a shelling.

The digit "0" or "1" at the i^{th} position of the binary string imposes the inequality $x_i \geq 0$ or $x_i \leq 0$, respectively, on the simplex described by the string. If we introduce a third symbol "$*$" to indicate that both inequalities are imposed simultaneously, or in other words $x_i = 0$, we now have a unique name for each of the 3^n simplices of all dimensions in the boundary complex of the cross polytope, including the empty simplex whose name is a string of all $*$'s. This notation and the rule $0 \cap 1 = *$ give us a purely symbolic way to find the intersection of any two simplices.

At each stage of a partial shelling, we have, in the case of a simplicial complex, only the requirement that every simplex after the first be attached along a non-empty union of its facets. (The maximal simplices are themselves facets, of dimension $n-1$, of the n-dimensional cross polytope, but the facets to which we are now referring are the faces of dimension $n-2$ of such an $(n-1)$-simplex, or in other words the simplices described by a string containing exactly one $*$.) Let S_k be a simplex which is proposed as an extension of a partial shelling $S_1, \ldots, S_{k-1}$. For each $j < k$, the set $\Delta_{j,k}$ lists all positions at which S_j and S_k differ (in other words, $\Delta_{j,k}$ is the set of $*$'s in $S_j \cap S_k$). The simplex $S_j \cap S_k$ is part of the set along which S_k is attached to the partial shelling, and the attachment set must be a union of facets. In other words, there must be some some $i \in \Delta_{j,k}$ and some $h < k$ such that $\Delta_{h,k} = \{i\}$, and this condition must hold for all $j < k$. If all of these conditions hold, the proposed extension of the partial shelling is legal.

It is easier to characterize these conditions in terms of the set of simplices *not* in the partial shelling. Let $\overline{\Delta}_k$ represent the set of positions i such that $\Delta_{h,k} \neq \{i\}$ for any $h < k$. Then S_k is a valid extension of the partial shelling as long as the inclusion $\Delta_{j,k} \subseteq \overline{\Delta}_k$ does not hold for any $j < k$. In other words, all of the $2^{|\overline{\Delta}_k|}$ simplices differing from S_k in some subset of $\overline{\Delta}_k$ must be in the complement of the partial shelling $S_1, \ldots, S_{k-1}$.

2.2.4 Pseudoplane arrangements

A *pseudoplane arrangement* of dimension d will be, for our purposes, a closed topological d-ball, called the *window* of the arrangement, together with a finite collection of properly and smoothly embedded $(d-1)$-balls, called *pseudoplanes*, such that the intersection of any k pseudoplanes is either empty or is a properly embedded ball of dimension exactly $d-k$. (Properly embedded means that the interior lies in the interior and boundary lies in the boundary; in the case of a properly embedded 0-ball, this means the point lies in the interior.) We allow also for any finite number of additional empty pseudoplanes not intersecting the window of the arrangement, and we allow for the case of an empty window, in which case the dimension of the pseudoplane arrangement is said to be -1. To the term "pseudoplane arrangement" we should perhaps add the qualifications "compact" (because the arrangement covers not all of space but only a finite window), "generic" (because the pseudoplanes are assumed to be in general position), and "affine" (because the pseudoplanes model affine hyperplanes rather than linear subspaces of $\mathbb{R}^n$), but all of the pseudoplane arrangements we consider will possess these qualities, and we will suppress the profusion of adjectives.

One source of pseudoplane arrangements of dimension $\leq d$ comes from taking a set of affine hyperplanes in $\mathbb{R}^d$ and intersecting them with any closed, convex region (the window), provided that the hyperplanes are in general position both with respect to each other and with respect to the window. Such a pseudoplane arrangement is called *straight*. It will be important, however, that not all pseudoplane arrangements are combinatorially equivalent to a straight arrangement.

Suppose we are given an arrangement of n pseudoplanes. Each pseudoplane in the list (including any empty ones) can be given a co-orientation which partitions the d-ball into three sets: the pseudoplane itself, labelled "$*$", and two halves labelled "0" (in the positive direction) and "1" (in the negative direction) with respect to that pseudoplane. This allows us to associate each open region of the pseudoplane arrangement with a binary string of length n, and hence with a facet of the n-dimensional cross polytope. In fact, every point of the arrangement lies in a region corresponding uniquely to some simplex of the boundary complex of the cross polytope. If the region has dimension $d - k$, the corresponding simplex in the boundary of the cross polytope will have dimension $n-k-1$, or in other words will have a label with "$*$" in exactly k positions.

2.3 Constructing the counterexample

The machinery of pseudoplane arrangements allows us to construct subcomplexes of the boundary complex of the cross polytope which are shellable, but whose shellings cannot be extended. The main result of this chapter is the following:

Theorem 2.3.1. *The cross polytope of dimension* 12 *or greater does not have an extendably shellable boundary complex.*

2.3.1 Pseudoplane arrangements as partial shellings

Every pseudoplane arrangement of interest to us will be interpreted, if possible, as the complement of a partial shelling of the cross polytope. In this setting, it is important to be able to recognize in what ways a partial shelling can be extended by a single facet. Which regions are *removable*? Which regions correspond to facets that can be added to the shelling, i.e. removed from the complement of the shelling?

If $\mathcal{A}$ is a pseudoplane arrangement, a compact subset Σ of $\mathcal{A}$ is called a *sector* of $\mathcal{A}$ if $\partial\Sigma$ is composed entirely of pieces of pseudoplanes and the boundary of the window, and if no pseudoplane intersects the interior of Σ. Each sector of $\mathcal{A}$ is associated with a unique binary string, and hence with some facet of the cross polytope. In the case of interest, the set of sectors of $\mathcal{A}$ corresponds exactly to the complement, in the cross polytope, of the set of simplices $S_1, \ldots, S_{k-1}$ of some partial shelling.

Lemma 2.3.1. *Let $\mathcal{A}$ be a d-dimensional arrangement of n pseudoplanes corresponding to the complement of a partial shelling $S_1, \ldots, S_{k-1}$ of the n-dimensional cross polytope, and let Σ be a sector of $\mathcal{A}$ corresponding to a simplex S_k of the cross polytope. Then $S_1, \ldots, S_k$ is a partial shelling if and only if the set of pseudoplanes adjacent to Σ has a non-trivial intersection.*

Proof. We have a characterization of whether S_k is a legal addition to the partial shelling in terms of binary strings: List all simplices *not* in the shelling which differ from S_k in a single entry, and let $\overline{\Delta}_k \subseteq \{1, \ldots, n\}$ list the indices for which this happens. Then S_k is a legal addition if and only if each of the $2^{|\overline{\Delta}_k|}$ simplices which differ from S_k in exactly some subset of the indices $\overline{\Delta}_k$ also lies outside of the partial shelling $S_1, \ldots, S_{k-1}$.

The set $\overline{\Delta}_k$ is exactly the set of pseudoplanes adjacent to Σ, and if there is a nontrivial intersection of all these pseudoplanes, then a neighborhood of that intersection is divided into $2^{|\overline{\Delta}_k|}$ regions, one for each simplex required to lie outside of the partial shelling $S_1, \ldots, S_{k-1}$. Since the sectors of $\mathcal{A}$ correspond exactly to the complement of the partial shelling, S_k is thus a legal addition.

Suppose on the other hand that each of the required $2^{|\overline{\Delta}_k|}$ sectors exists in $\mathcal{A}$, and in particular that the sector $\overline{\Sigma}$ exists whose binary label differs from that of Σ in exactly the entries of $\overline{\Delta}_k$. Then $\Sigma \cap \overline{\Sigma}$ is the non-empty intersection of all the pseudoplanes of $\overline{\Delta}_k$. □

Remark 2.3.2. It is also true, although we do not need it for our results, that a sector Σ is "removable" if and only if the property of being a pseudoplane arrangement is preserved when an open regular neighborhood of Σ is excised. It follows that every partial shelling corresponds to the complement of some pseudoplane arrangement, since the set of coordinate hyperplanes in $\mathbb{R}^n$, intersected with a closed ball containing the origin, is a pseudoplane arrangement corresponding to the complement of the empty partial shelling, of which every partial shelling is an extension.

2.3.2 Dissolving and undissolving straight arrangements

If it were not for the existence of non-straight pseudoplane arrangements, the cross polytope would be extendably shellable in all dimensions, as the following lemma demonstrates.

Lemma 2.3.2. *Let $\mathcal{A}$ be a pseudoplane arrangement of dimension d which is straight, coming from the intersection of a convex, compact window $W \subset \mathbb{R}^n$ with n affine hyperplanes. Then there is a partial shelling of the n-dimensional cross polytope whose complement corresponds to the sectors of $\mathcal{A}$, and the same partial shelling can be extended to a complete shelling.*

Proof. We may assume without loss of generality that $d = n$, by embedding the given system in $\mathbb{R}^n$ and extending the hyperplanes in any consistent manner. We may also assume that W is full-dimensional, by taking a small neighborhood, and that the origin of $\mathbb{R}^n$ lies in the interior of W. Finally, we will make strong assumptions on the general position of the hyperplanes, the first of which is that no hyperplane passes through the origin and that they intersect in a unique point $p \in \mathbb{R}^n$.

Consider the family of straight pseudoplane arrangements given by taking tW as the window instead of W, where t is any positive real number. For large enough values of t, tW will contain the point p, and these arrangements will have a full complement of 2^n sectors. For small enough values of t, on the other hand, the corresponding pseudoplane arrangement will consist of a single sector in the neighborhood of the origin, disjoint from all hyperplanes. The strong assumption of general position we make is this: that for every integer $1 \leq k \leq 2^n$, there exists some value of t for which the arrangement tW has exactly k sectors. In other words, we require that no two sectors happen to "disappear" at precisely the same moment as t decreases towards zero.

The well-defined order of disappearance so obtained, concluding with the sector of the origin, is a shelling order on the facets of the cross polytope in dimension n. In particular, the sectors of $\mathcal{A}$ correspond to the complement of that stage of the partial shelling for which $t = 1$, which shelling can be completed. $\square$

The converse is false: there exist pseudoplane arrangements which cannot be straightened but which nonetheless correspond to extendable partial shellings. In particular, it is a somewhat subtle exercise to show that every pseudoplane arrangement of dimension 2 (not all of these can be straightened!) has a removable region.

2.3.3 An obstruction in dimension 3

This brings us to a discussion of pseudoplane arrangements of dimension 3, where a phenomenon related to the existence of knots in this dimension will lead us to a construction of non-extendable partial shellings for $d = 3$ (and $n = 12$). We do this by constructing (and shelling as far as) a pseudoplane arrangement of dimension 3 without a single removable region. The regions we must eliminate, touching the boundary of the 3-ball, can take one of three possible forms:

1. A region enclosed by a single pseudoplane intersecting no others and tracing out a circle on the boundary;
2. A region enclosed by two pseudoplanes intersecting along an arc and tracing out a bigon on the boundary whose vertices are the endpoints of that arc; or
3. A region enclosed by three pseudoplanes forming, with the boundary, a topological tetrahedron.

Naturally the third case is the most interesting, and most difficult to eliminate.

Instead of single pseudoplanes, we will work with pairs of nearly-parallel pseudoplanes intersecting along an arc which remains close to the boundary of the 3-ball. Ignoring the pseudoplanes themselves for the moment, and keeping track of just the arcs of intersection, we can encode the most important information with just a picture of those how those arcs project onto the 2-sphere boundary. In other words, we draw diagrams in the plane of arcs which cross over and under each other. The result is much like a diagram of a knot or link, but one in which every curve has endpoints. A *weaving diagram* is a collection of arcs in the plane (considered as part of the 2-sphere), together with information at each intersection of which arc crosses "over" and which arc crosses "under". We require that the arcs and their endpoints be in general position, and furthermore that no arc cross itself, and that any given pair of arcs cross no more than once.

A weaving diagram with k arcs may encode a three-dimensional arrangement of $2k$ pseudoplanes, although the correspondence is not uniquely specified; some choices need to be made about how pairs of pseudoplanes intersect with other pairs of pseudoplanes in the interior of the 3-ball, or near the boundary but on the opposite side from where the pseudoplanes of a pair cross each other. (Every arc is actually half of a loop, but the other half is not drawn.)

Each endpoint of an arc in a weaving diagram, if it is traced back the last place at which the arc crossed another, gives rise to a triangle on the boundary of the pseudoplane arrangement. It is to these triangles most particularly that we must pay attention as we attempt to constuct a pseudoplane arrangement without removable regions. The rule is simple: as long as the last crossing before the endpoint is an undercrossing, the region of the 3-ball which includes this triangular region of the 2-sphere will be bounded by at least 4 pseudoplanes rather than 3, and hence will not be removable, will not give rise to a legal addition to a shelling of the cross polytope.

Figure 2.1 depicts a region near the boundary of a pseudoplane arrangement where two arcs cross before hitting the boundary. The two pairs of planes are also depicted to show how the endpoint of the "untucked" arc (the vertical arc, in the example) gives rise to a removable region, while the (horizontal) "tucked" one does not.

The first goal, then, is to construct a weaving diagram in which every endpoint is "tucked in". (Such a diagram could be realized physically by weaving sticks around a ball; the endpoint criterion is exactly what is required to ensure that the woven structure

Figure 2.1: Only “untucked” endpoints are removable

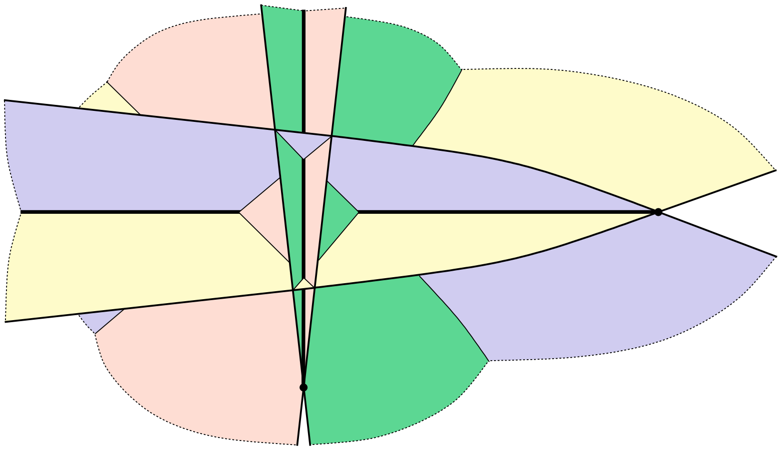

Figure 2.2: A weaving diagram with all endpoints tucked.

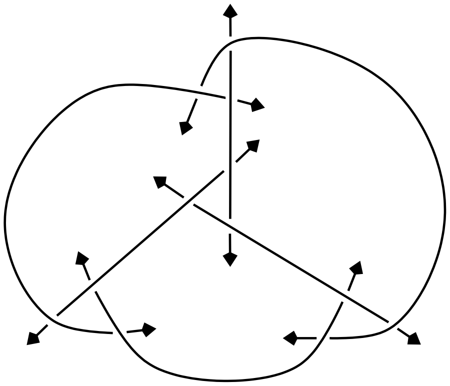

does not fall apart.) Such a construction is easy enough with a sufficent number of arcs; let us try to find a minimal example. Each arc, in addition to having its own endpoints held in place, must in turn secure the endpoints of at least two other arcs on average; it follows that the average number of other arcs crossed by any given arc must be at least 4, and hence that the total number of arcs must be at least 5. With exactly 5 arcs, however, every arc must cross every other arc exactly once, and the fact that the complete graph on 5 vertices is non-planar prevents this. Allowing ourselves a total of 6 arcs leads quickly to the weaving diagram pictured in Figure 2.2. This solution is unique up to a choice of clockwisdom for each of the four triangular regions. (We have not made the most symmetric choice of orientations, but rather a choice which simplifies the realization of this weaving diagram as a pseudoplane arrangement.)

2.3.4 Weaving an arrangement

In order to prove the main claim, it remains to

1. Construct a pseudoplane arrangement corresponding to the above weaving diagram,
2. Verify that none of the boundary regions of this arrangement is removable, and
3. Show that this arrangement comes from a partial shelling.

We start with a specific one-parameter family of planes in $\mathbb{R}^3$, linear in the parameter. (Equivalently, we start with a specific arrangement of hyperplanes in $\mathbb{R}^4$.) For the topological 3-ball we take the unit cube in Cartesian 3-space bounded by the six equations $x = \pm 1, y = \pm 1, z = \pm 1$; the parameter t will also take values between -1 and 1.

The twelve planes (or rather families of planes, or rather twelve hyperplanes in (x, y, z, t)-space) will be perturbations of the six pairwise diagonals $y = \pm x$, $z = \pm x$, and $z = \pm y$. The desired diagram will appear, after perturbation and resolution to a single, non-straightenable, 3-dimensional arrangement of pseudoplanes, in a neighborhood of the edges of the tetrahedron

$$\begin{aligned} x + y + z &\leq 1, \\ x - y - z &\leq 1, \\ -x + y - z &\leq 1, \\ -x - y + z &\leq 1 \end{aligned}$$

or in other words along the union of the sets $\{y = x, z = -1\}$, $\{y = -x, z = 1\}$, $\{z = x, y = -1\}$, $\{z = -x, y = 1\}$, $\{z = y, x = -1\}$, and $\{z = -y, x = 1\}$.

Starting with two copies of each of the six exact diagonal planes, we apply a series of perturbations until we obtain a family of planes which are in general position for all values of t outside some neighborhood of 0. The perturbations are of successively smaller orders so that no perturbation affects the combinatorics of any resolutions already made of the initial failure of general position.

In order, then:

1. Rotate the planes $x = \pm y$ clockwise around the z-axis. (This gives, from the edges of the already-specified tetrahedron, the choice of orientation in the triangles of the target weaving diagram.) Displacement distance at the boundary: ε^3.

2. Raise (uniformly increase the z-coordinates of every point in) the planes $z = \pm x$ and lower the planes $z = \pm y$. Displacement distance: ε^4.

3. Further raise the pair of planes $z = x$ and further lower the pair of planes $z = y$. Displacement distance: ε^5. Up to this point, the planes still exist in identical pairs.

4. For each of the six pairs of planes, tilt the two planes slightly in opposite directions around a line which is just below the surface of the unit cube, a line close to one of the six specified edges of the tetrahedron. Maximum distance between planes at the opposite side of the cube: ε^6.

   Each pair of planes should intersect at exactly the same depth ε below the surface of the cube. (This is an exception to the use of successively smaller constants, but the actual displacement of any one plane is still of order ε^6, so the combinatorics of either of the two planes with respect to all other planes is unaffected.)

5. It is at this point that we first make use of the parameter $-1 \leq t \leq 1$. We choose six different constants $p_1, \ldots, p_6$, one for each of the six pairs of planes. For each pair of planes, make a linear adjustment, as a function of t, which displaces no single plane a distance of order more than ε^6, but in such a way that the depth at which pair i intersects, rather than exactly ε, is $\varepsilon + tp_i\varepsilon^2$. Since no two of the constants are the same, we can examine any two distinct pairs of planes i, j, and one of the pairs (say pair i) will intersect at a deeper level than pair j at $t = -1$, but the order will be reversed (pair j deeper than pair i) at $t = 1$. This will eventually give us independent control, at each of 12 crossings, of which of two arcs gets woven below the other.

The result is a straight hyperplane arrangement in general (enough) position in $\mathbb{R}^4$, and it follows from Lemma 2.3.2 that there is a partial shelling of the 12-dimensional cross polytope whose complement corresponds to the regions of this arrangement. The 4-dimensional hyperplane arrangement is, in most places, simply the product of a 3-dimensional arrangement with the interval $-1 \leq t \leq 1$; the 12 localized places where this is not the case correspond exactly to the 12 crossings of the weaving diagram. At each of the 12 exceptional places, there is one removable region of the 4-dimensional arrangement at the $t = -1$ boundary of the hypercube, and another removable region at the boundary $t = 1$. (From the 3-dimensional perspective, the regions affected are entirely within the interior of the cube although very close to the boundary; from the perspective of the fourth dimension t, these regions touch the boundary and extend, in a tiny sliver, all the way to some depth close to $t = 0$.)

At each of these 12 places, we are allowed to remove one of the two regions, but not both, and the result is a trivialization in the t direction, where the remaining 3 dimensions have been resolved either to an over- or an under-crossing in the weaving diagram. Most importantly, the dozen choices can be made entirely independently; no one choice affects the combinatorics at any of the other twelve sites.

We extend the partial shelling given to us by Lemma 2.3.2 with 12 additional facets of the cross polytope, by trivializing in the t direction at each of the 12 sites in such a way that the six lines of intersection between pairs of pseudoplanes agree with the under- and over-crossings of the weaving diagram. As a result we obtain a non-straightenable pseudoplane arrangement which is now homeomorphic to the product of a 3-dimensional arrangement with an interval, and which is thus combinatorially equivalent to a 3-dimensional arrangement corresponding to the desired weaving diagram. The most troublesome triangles in the boundary of this arrangement are taken care of—we have ensured that none of these gives rise to a removable region—and a bit of checking verifies that no other triangle in the boundary is removable. It follows that there is no facet of the 12-dimensional cross polytope with which to extend this partial shelling, and the 12-dimensional cross polytope is not extendably shellable.

For this particular construction, a minimum of 12 pseudoplanes were required; it remains an open question in what dimension the cross polytope becomes non-extendably shellable.

Note that the wrong choice in any of the 12 pairs of removable regions leads to an extendable shelling; this is not an example that one would expect to find easily in a random search. Indeed, attempts to find a non-extendable shelling by randomized computer search, even intelligent searches which seek to minimize the number of future choices, have so far met with no success: in some imprecisely defined way, the cross polytope seems to be "nearly always" extendably shellable, in any dimension low enough that it is accessible to search by computer.

2.4 A uniform oriented matroid with a mutation-free element

The 3-dimensional pseudoplane arrangement we have constructed happens to satisfy properties that ensure that it is in fact a pseudosphere arrangement: Every pair or triple of planes has a non-trivial intersection (an arc or a point, repectively) and so the boundary has the right symmetry to be considered a copy of the real projective plane rather than a 2-sphere, one of 13 topological RP^2s in a topological RP^3. By a well-known correspondance between pseudosphere arrangements and uniform oriented matriods, and reinterpreting the meaning of a “removable region”, this gives us a smaller example than was previously known (13 instead of 17 [BR01] or 21 [RG93]) of a uniform oriented matroid with a mutation-free element.

2.5 Verification by computer

Given such a concrete counterexample, it is not too difficult to verify it directly by computer. The first three perturbations of the arrangement of six diagonal planes were done explicitly, and a labelling obtained for each of the complementary regions. The planes were then doubled, and additional regions corresponding to the weaving diagram (including the 12 crucial choices) were calculated by hand. The result was a list of 299 facets, each encoded as a binary string of length 12.

A few attempts were made to shell all but the 299 facets using various systematic greedy algorithms, all of which failed, but randomized attempts at a shelling of the 3797-simplex subcomplex succeed with reasonably high probability, and in these cases the computer is able to verify that none of the 299 remaining facets is a legal next move. (The 3797-simplex subcomplex is itself shellable but not extendably shellable. To add to the confusion, in every case encountered, when a partial shelling of the subcomplex could not be extended to a complete shelling of the subcomplex, it could nonetheless be extended to a shelling of the cross polytope once the 299 “forbidden” facets were allowed—the particular set of complementary facets appears to be quite delicate, and no search to date by computer has produced an acceptable alternative.) One such non-extendable shelling, and a list of the 299 facets in the complement, is available at

http://math.berkeley.edu/~hthall/stuck.html

Frank Lutz has verified the validity of this partial shelling, as well as its non-extendability, using software he developed for use with GAP.

Chapter 3

Chamber Complexes of High Complexity

In this chapter we consider a separate problem, still involving polytopes and their subdivisions, but this time in only two dimensions. As before, the object we consider has a straightforward and uncomplicated construction, but, also as before, a thorough analysis of sufficiently large examples reveals a richer structure than had been thought possible to exist. In the case of the cross polytope, the essential complexity, although ostensibly in 12 dimensions, comes actually as a result of the phenomenon of knotting in 3 dimensions. In the case we next consider, the crucial principle at play seems to be this: Rational numbers are poor approximations of other rational numbers.

A convex polygon in the Cartesian plane is called a lattice polygon if each of its vertices is a lattice point, i.e. has integer coördinates. For every pair of lattice points anywhere in the polygon we draw the line segment connecting them; the resulting subdivision of the lattice polygon into (much) smaller polygonal chambers is called the chamber complex. We define the complexity of the original lattice polygon as the maximal number of edges of any single chamber. In small examples, most or all of the chambers appear to have 3 or 4 edges. A simple Euler characteristic calculation reveals that the average number of edges can be no greater than 4, but we can gain additional insight into why most or nearly all chambers should have few edges by considering what has to happen for the number of edges of a particular chamber to increase as the surrounding lattice polygon is expanded to include more lattice points.

Given a chamber which already has a certain number of edges, and whose number of edges we wish to increase, we firstly must avoid cutting through the middle of the chamber, which would reduce the number of edges of the resulting pieces by nearly a factor of two. For any additional cut to increase the number of edges, it must cut off a single vertex of the existing chamber. More to the point, however, it must pass between that vertex and its two neighbors, without crossing exactly through any of the three. Here is where the poor approximation of rationals by rationals comes into play. Specifically, the three vertices between which the new line must cross are rational points of relatively low denominator. In order for the new line to miss all of them, it must create two new nearby intersections, also of low denominator; in most cases it is more likely for the new line to pass through at least one of the three. In order for the number of edges of a chamber to increase without bound, this unlikely process must happen again and again, without the chamber ever being split by a stray line.

The complexity of lattice polygons is important because of its relation to a 1997 conjecture [ST97] (the so-called "3/4 conjecture") of Sturmfels and Thomas. This question arose in the context of complexity bounds on certain calculations involving Gröbner bases, and the conjecture was well supported by all calculations that they were able to carry out at that time and in the original context of the problem. Subsequent work, however, including a theorem [HMS04] of Hoşten, Maclagan, and Sturmfels, has reduced the conjecture to the following question: Is there a uniform bound on the complexity of lattice polygons? The stronger form of the conjecture fails already with the existence of pentagons within the chamber complex of lattice polygons, and the entire conjecture is refuted if it can be shown that lattice polygons exist of arbitrarily high complexity.

Computerized searches have located chambers with as many as 15 vertices, but no clear pattern has emerged from these examples, and the question of a uniform bound remains open. The present work exhibits a sequence P_n of lattice polygons and verifies, for each $n \geq 1$, that the chamber complex of P_n contains a polygon with at least $4n$ sides.

The success of this construction actually relies on the same principle that would seem to prevent it: the poor approximation of rational numbers by other rational numbers of low denominator. The key is to find some point whose coördinates are rational numbers of low denominator, but through which no line can *ever* pass as long as the lattice polygon is constrained in one of its dimensions. If we then increase the size of the lattice polygon in the remaining dimension, and are careful about the manner in which we do so, the lines which are forced to approximate (poorly) the forbidden point may approach it from all sides at similar rates, creating a chamber with many edges.

3.1 Definitions

3.1.1 Lattice Polygons

We work in the Cartesian plane $\mathbb{R}^2$, with the horizontal and vertical axes labelled x and y, respectively. When x and y are both integers, (x, y) is called a *lattice point.* A *lattice polygon* is a convex polygon each of whose vertices is a lattice point. Equivalently, a lattice polygon is the convex hull of a finite non-collinear set of lattice points.

3.1.2 The Chamber Complex

The main object of study is a construction that subdivides a lattice polygon along certain line segments. If P is a lattice polygon and $(x_1, y_1), (x_2, y_2)$ are lattice points of P—points either on the boundary of P or in the interior—then the line segment from (x_1, y_1) to (x_2, y_2) is called a *fissure* of P. The *chamber complex* of P, denoted $\mathscr{C}(P)$, is a collection of smaller polygons, called the *chambers* of P, obtained by splitting P simultaneously along all of its fissures. To be precise, $\mathscr{C}(P)$ consists of those polygons C such that each edge of C is part of a fissure of P, but no fissure of P passes through the interior of C. Each vertex of a chamber is either a lattice point or an intersection of fissures, from which it follows that chamber vertices are rational points. It is also true that chambers are convex: Suppose some chamber had a reëntrant vertex v. It could not be an intersection of fissures or a point on the boundary of P, and in the remaining case where v is an interior lattice point, the fissures emanating from v to the vertices of P form a fan whose angles are all less than π.

To each lattice polygon P we associate a positive integer, the *complexity* of P, defined as the maximum over all of $\mathscr{C}(P)$ of the number edges of a chamber. Our main result is a family of polygons over which the complexity is unbounded.

3.1.3 Flush-left Polygons and Barriers

We use the term *flush-left polygon* to refer to a lattice polygon that is "left-justified" along the y-axis. To be precise, P is a flush-left polygon if every vertex of P has non-negative x-coördinate, with exactly two of the vertices lying along the y-axis. If the maximum x-coördinate of any vertex is w, then w is called the *width* of the flush-left polygon. Every lattice polygon is equivalent to a flush-left polygon via some lattice-preserving affine transformation, but the width will depend on which edge is sent to the y-axis.

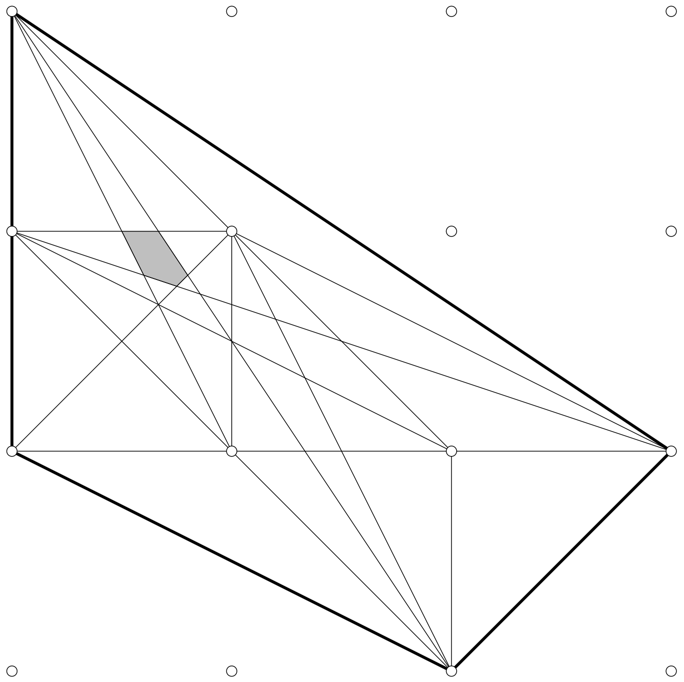

Figure 3.1: A flush-left polygon of complexity 5

Example 3.1.3.1. Figure 1 depicts a flush-left polygon of width 3 and its associated chamber complex. Among the chambers in Figure 1 we count 26 triangles, 5 quadrilaterals, and a single (shaded) pentagon, so this is an example of a lattice polygon of complexity 5.

Although lattice polygons are determined by a discrete set of vertices, it will be convenient to regard them as approximations of continuous objects. For w a positive integer, we call a function $u : [1, w] \to \mathbb{R}$ a *barrier of width* w if u is positive and concave downward and satisfies the technical assumptions

$$u(1) \geq u(w) \tag{3.1}$$

and

$$u(x) \leq \frac{u(1)}{2}(x+1) \quad \text{for all } x \in [1, w]. \tag{3.2}$$

We make no distinction between a barrier and its graph. Property 3.2 says that a barrier u lies entirely on or below the line that passes through the points $(-1, 0)$ and $(1, u(1))$. Given concavity and assuming $w > 1$, the derivative of a barrier is defined (if finite) at

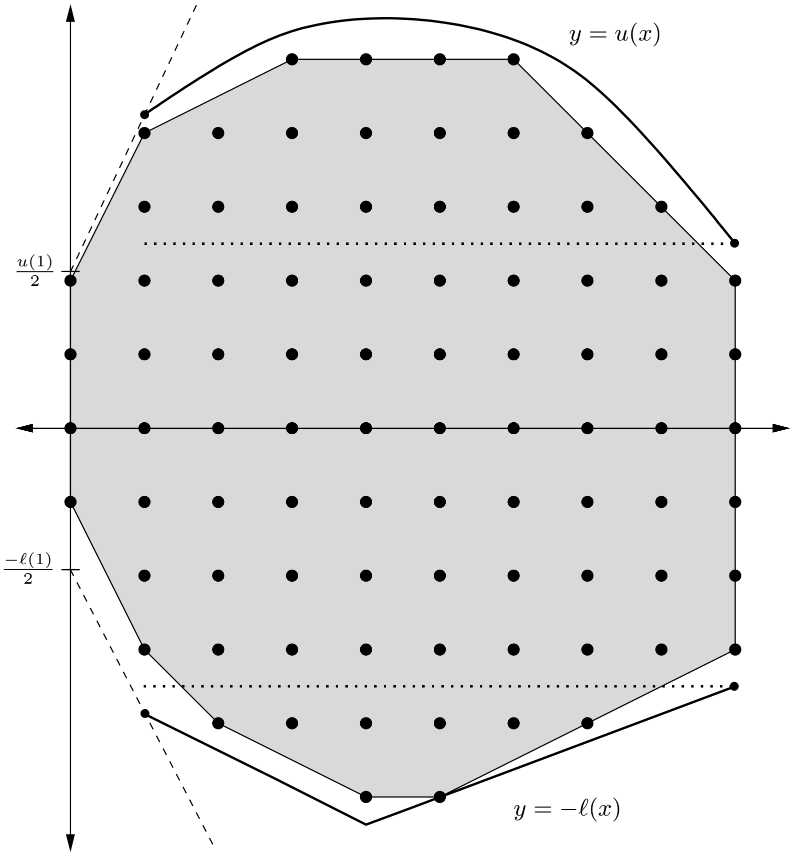

Figure 3.2: The polygon with upper barrier u and lower barrier ℓ

the left endpoint, and Property 3.2 under these assumptions is equivalent to

$$u'(1) \leq \frac{1}{2}u(1). \tag{3.3}$$

Every secant line of a barrier u has positive y-intercept, and as x increases, the slope of the line between the origin and the point $(x, u(x))$ does not increase. In other words,

$$\frac{u(x_1)}{x_1} \geq \frac{u(x_2)}{x_2} \quad \text{whenever } 1 \leq x_1 \leq x_2 \leq w. \tag{3.4}$$

We will use this inequality in the proof of the Fissure Slopes Lemma at the end of Section 3.5.

Taking one barrier as an upper bound and the negative of another barrier as a lower bound, where both barriers are of the same width w, we obtain a flush-left polygon of width w once we have also specified upper and lower bounds on the y-axis, where the barriers are not defined. If u and ℓ are barriers of width w, the polygon *with upper* *barrier* u *and lower barrier* ℓ is defined to be the convex hull of those lattice points (x, y) satisfying one of

$$1 \leq x \leq w, \quad -\ell(x) \leq y \leq u(x)$$

or

$$x = 0, \quad \frac{-\ell(1)}{2} \leq y \leq \frac{u(1)}{2}.$$

Concavity and Property 3.2 ensure that this set of lattice points already lies within a convex region, so taking the convex hull of these points does not introduce any new lattice points. The polygon so produced is a flush-left polygon of width w—except in the degenerate case where only one lattice point lies on the y-axis, which is avoided if $u(w) \geq 2$ or $\ell(w) \geq 2$. The use of a distinct bound for the case $x = 0$ may seem somewhat arbitrary (since we could have simply defined barriers over the interval $[0, w]$ in the first place) but the fissures of most interest to us are those with exactly one endpoint on the y-axis, and the particular bounds given for the case $x = 0$ ensure (as we will demonstrate in the proof of the Fissure Slopes Lemma) that there are enough left-hand endpoints to appropriately complement the full range of choices for right-hand endpoints.

Example 3.1.3.2. Figure 2 shows a barrier u, the negative of a barrier ℓ, and the polygon with upper barrier u and lower barrier ℓ. Dotted lines illustrate compliance with Properties 3.1 and 3.2.

Property 3.1 provides that $u(w)$ is a uniform lower bound on $u(x)$ for $1 \leq x \leq w$, which is required at one point in the proof of the Fissure Slopes Lemma. Except for positivity, the conditions on barriers are homogeneous, so any positive multiple of a barrier is still a barrier. For example, we can avoid the degenerate case of one vertex on the y-axis by taking a large enough multiple to ensure that $u(w) \geq 2$.

We will often use the same function u to delineate both the upper and lower barriers of a polygon, in which case we refer simply to the polygon *with barrier* u.

3.1.4 Pseudofissures and pseudochambers

Just as we view a lattice polygon as a discrete approximation of a continuously-defined region, the fissures of most interest to us approximate certain lines whose slopes are not necessarily rational. These lines, called *pseudofissures,* are defined by a function, a point, and an index. Given a function $u : [1, w] \to \mathbb{R}$, a point (x_0, y_0) with $0 < x_0 < 1$, and an integer j in the range $1 \leq j \leq w$, the j^{th} *pseudofissure of* u *through* (x_0, y_0) is the line that passes through the two points

$$(x_0, y_0) \text{ and } (j, u(j) + y_0).$$

We are interested particularly in pseudofissures of a barrier or the negative of a barrier: If u and ℓ are barriers and P is the polygon with upper barrier u and lower barrier ℓ, then pseudofissures of u and of $-\ell$ may be approximated by honest fissures of P. More specifically, we will see that in many cases there is a family of fissures all starting on the y-axis, all ending at the j^{th} column of lattice points, and all crossing through a common point (x_0, y_0), such that the fissures of maximum and minimum slope within the family are approximations, respectively, of the j^{th} pseudofissure of u and j^{th} pseudofissure of $-\ell$ through (x_0, y_0). This is a reasonable expectation, as long as there is no obstruction to fissures ending at the j^{th} column and passing through the point (x_0, y_0), because the pseudofissures pass through points that are near the upper and lower bounds which delimit that column of lattice points. Rather than defining pseudofissures to intersect the line $x = j$ at precisely the points $(j, u(j))$ and $(j, -\ell(j))$, we have opted for a definition which gives the j^{th} pseudofissure of u a slope of exactly

$$\frac{u(j)}{j - x_0} \tag{3.5}$$

regardless of the value of y_0.

A finite collection of pseudofissures cuts the plane into regions, and such a region, when bounded, is called a *pseudochamber*. To be precise, given a finite collection Ψ of pseudofissures, a polygon D is a Ψ-pseudochamber if each edge of D lies on a pseudofissure in the collection Ψ, but no pseudofissure of Ψ passes through the interior of D. One of our essential results (Lemma 3.2.3) is that certain pseudochambers arising from a barrier u and the negative of a barrier ℓ are approximated (to arbitrary relative precision) by honest chambers of the lattice polygon with upper barrier u and lower barrier ℓ.

3.2 Lemmata

In order to construct chambers of any size and shape with the required degree of precision, we require a few preliminary results giving us a fine degree of control over the placement and slopes of fissures within a lattice polygon defined by a pair of barriers.

3.2.1 Placement of fissures

The first step towards knowing where to find the promised $4n$-gon is the following lemma, which proves the existence of points through which fissures can never cross exactly, but can rather only approximate (ideally from all sides). More precisely, it tells us, for certain values of x_0 between 0 and 1, the permissible values of y_0 for which there may exist fissures crossing through the point (x_0, y_0). The statement of the lemma makes use of the following definition: Given a positive integer k, the *odd part* of k is the unique odd integer i such that $k = 2^\alpha i$ for some integer α.

Lemma 3.2.1. *Let P be a flush-left polygon of width w. Choose a positive integer β and an odd positive integer q such that $w < 2^\beta < q$, and set $x_0 = \frac{2^\beta}{q}$. Suppose (x_0, y_0) lies on a fissure F of P, that k is the x-coördinate of the right-hand endpoint of F, and that i is the odd part of k. Then y_0 is a multiple of $\frac{1}{iq}$.*

Proof. Let P, w, β, q, and x_0 be as in the statement of the lemma, and suppose that a fissure F of P passes through the point (x_0, y_0) for some y_0. Since $0 < x_0 < 1$, exactly one endpoint of F must lie on the y-axis. The endpoints of F are thus $(0, a)$ and (k, b) for some integers a, b, and $1 \leq k \leq w$. We factor out the odd part of k as $k = 2^\alpha i$, noting that this implies $\alpha < \beta$. The fissure F is a segment of the line whose equation is

$$y = \frac{b-a}{k}x + a,$$

and in particular

$$y_0 = \frac{b-a}{2^\alpha i}\frac{2^\beta}{q} + a = \frac{(b-a)2^{\beta-\alpha} + aiq}{iq}, \tag{3.6}$$

which completes the proof. □

3.2.2 Slopes of fissures

This lemma has a converse of sorts: For a certain class of points (x_0, y_j) and for suitable barriers u and ℓ, we can guarantee, for the polygon with upper barrier u and lower barrier ℓ, that some fissure crosses through the point (x_0, y_j). What is more, we can determine, to within a certain tolerance, the slopes of all such fissures.

Lemma 3.2.2 (Fissure Slopes). *For any positive integer w, choose a positive integer β and an odd positive integer q such that $w < 2^\beta < \frac{1}{4}q$, and set $x_0 = \frac{2^\beta}{q}$. Choose j an odd integer with $1 \leq j \leq w$, and let y_j be a multiple of $\frac{1}{jq}$ such that y_j is not a multiple of $\frac{1}{iq}$ for any odd positive integer $i < j$. Let u and ℓ be barriers of width w satisfying*

$u(1), \ell(1) \geq 16|y_j|$; $u(w), \ell(w) \geq wq - 2^\beta$; and $\frac{5}{6} \leq \frac{u(1)}{\ell(1)} \leq \frac{6}{5}$. Let P be the polygon with upper barrier u and lower barrier ℓ. Denote by M the set of all slopes of fissures of P that pass through the point (x_0, y_j). Then

1. *M is non-empty: some fissure of P passes through (x_0, y_j),*
2. *The maximum value of M is within $q + \frac{4}{3}|y_j|$ of $\frac{u(j)}{j-x_0}$, and*
3. *The minimum value of M is within $q + \frac{4}{3}|y_j|$ of $\frac{-\ell(j)}{j-x_0}$.*

Lemma 3.2.2 is similar to Lemma 3.2.1, but provides more precise information about the slopes of fissures through the allowed points. This precision requires additional assumptions: the point of interest is now constrained to lie within the region $0 < x_0 < \frac{1}{4}$ near the left boundary of the lattice polygon; the denominator of y_j must take a specific form; and the barriers u and ℓ must be "large enough" with respect to certain fixed data and "close enough" to each other at their left endpoints. This greater degree of precision about fissure slopes also comes at the cost of a considerably longer proof, and we will defer the discussion of those details until after we have established the main results which depend on the lemma.

It is important that the quantity $q + \frac{4}{3}|y_j|$ is independent of the choice of barriers u and ℓ. In particular, the error bound remains constant when u and ℓ are replaced by large multiples of themselves, which means that we can make the relative error as small as desired simply by scaling the barriers.

Note also that the lemma provides us with a different pair of slope bounds for each odd integer in the range $1 \leq j \leq w$. It is this precise and independent control over slopes of fissures through different families of points (x_0, y_j), as j varies, that will allow us to construct a $4n$-gon around a single point (x_0, y_0)—a point which no fissure may traverse exactly, but which different families of fissures may approach at essentially independent rates.

3.2.3 Approximating pseudochambers

The Fissure Slopes Lemma gives us a sufficient degree of control over the minimum and maximum slopes of fissures through certain carefully chosen points that we can construct entire chambers to within an arbitrary degree of approximation.

Lemma 3.2.3 (Pseudochamber Approximation). *For any positive integer w, choose a positive integer β and an odd positive integer q such that $w < 2^\beta < \frac{1}{4}q$, and set*

$x_0 = \frac{2^\beta}{q}$. Let y_0 be any odd multiple of $\frac{1}{2q}$. Let u and ℓ be barriers of width w such that $\frac{5}{6} \le \frac{u(1)}{\ell(1)} \le \frac{6}{5}$. For each positive N such that $Nu(w) \ge 2$, denote by $P(N)$ the polygon with upper barrier Nu and lower barrier $N\ell$, and, for all N such that (x_0, y_0) lies within $P(N)$, let $C(N)$ be the chamber of $P(N)$ that contains that point. (By Lemma 3.2.1, no fissure contains the point (x_0, y_0), so the point lies in the interior of a unique chamber.) Stretch each such polygon $C(N)$ by a factor of N in the x-direction and call the result $\widehat{C}(N)$, or, more precisely, let $\widehat{C}(N)$ be the image of $C(N)$ under the affine transformation

$$(x, y) \mapsto (x_0 + N(x - x_0), y).$$

Now construct a collection Ψ of pseudofissures, consisting of four pseudofissures for each odd integer j in the range $1 \le j \le w$: the j^{th} pseudofissures of u through $(x_0, y_0 + \frac{1}{2jq})$ and, separately, through $(x_0, y_0 - \frac{1}{2jq})$, and the pair of j^{th} pseudofissures of $-\ell$ through the same two points. Let D represent the Ψ-pseudochamber containing the point (x_0, y_0). Then

$$\lim_{N\to\infty} \widehat{C}(N) = D.$$

In particular, for some N, $P(N)$ has a chamber $C(N)$ with at least as many edges as D.

The proof of this lemma will also be deferred until Section 3.5.

3.3 Chambers of any size

Assuming the validity of Lemma 3.2.2 and Lemma 3.2.3, we are now in a position to state and prove the main result of the chapter.

Theorem 3.3.1. *Given any positive integer n, there exists a flush-left polygon of width less than $2n$ whose complexity is at least $4n$.*

Proof of Theorem 3.3.1. Given any positive integer n, we exhibit a flush-left lattice polygon P of width $2n - 1$ and complexity at least $4n$.

Let β be an integer such that $2n \le 2^\beta$ (for example, $\beta = \lceil \log_2 2n \rceil$), let q be any odd integer such that $2^\beta < \frac{1}{4}q$ (for example, $q = 2^{\beta+2} + 1$), and let $x_0 = \frac{2^\beta}{q}$. Let y_0 be any odd multiple of $\frac{1}{2q}$ (for example, $y_0 = \frac{1}{2q}$). Define $u : [1, 2n-1] \to \mathbb{R}$ by

$$u(x) = \frac{x - x_0}{x}\sqrt{4n^2 - x^2}.$$

We will use u (or rather, a large multiple of it) for both the upper and the lower barriers of our lattice polygon; the assumption $\frac{5}{6} \le \frac{u(1)}{u(1)} \le \frac{6}{5}$ is satisfied trivially.

Construct a collection Ψ of pseudofissures consisting of four pseudofissures for each odd integer j in the range $1 \leq j \leq 2n-1$: the j^{th} pseudofissures of u through $(x_0, y_0 + \frac{1}{2jq})$ and, separately, through $(x_0, y_0 - \frac{1}{2jq})$, and the pair of j^{th} pseudofissures of $-u$ through the same two points.

We claim that the $4n$ pseudofissures in the collection Ψ are tangent to the circle

$$(x-x_0)^2 + (y-y_0)^2 = \left(\frac{1}{4nq}\right)^2$$

at the $4n$ points

$$\left(x_0 \pm \frac{1}{4nq}\sqrt{1-\left(\frac{j}{2n}\right)^2}, y_0 \pm \frac{1}{4nq}\left(\frac{j}{2n}\right)\right).$$

By symmetry, it is enough to verify this claim in the second quadrant of the circle. To this end, these three facts suffice:

1. Referring to Equation 3.5, the slope of the j^{th} pseudofissure of u through $(x_0, y_0 + \frac{1}{2jq})$ is
$$\frac{u(j)}{j-x_0} = \frac{1}{j-x_0}\frac{j-x_0}{j}\sqrt{4n^2-j^2} = \sqrt{\left(\frac{2n}{j}\right)^2 - 1}.$$
2. The point
$$\left(x_0 - \frac{1}{4nq}\sqrt{1-\left(\frac{j}{2n}\right)^2}, y_0 + \frac{1}{4nq}\left(\frac{j}{2n}\right)\right)$$
satisfies both the equation of the circle and the equation of the pseudofissure.
3. The line that contains both the center of the circle and the point of intersection has a slope which is the negative reciprocal of the slope of the pseudofissure.

It follows that the $4n$ pseudofissures are each tangent to the same circle at different points, and hence that the Ψ-pseudochamber D containing the point (x_0, y_0) is a $4n$-gon. If u is a barrier, then by the Pseudochamber Approximation Lemma, for some sufficiently large N the polygon P with barrier Nu will have a chamber with at least $4n$ edges, containing the point (x_0, y_0).

It remains only to prove, then, that

$$\frac{x-x_0}{x}\sqrt{4n^2-x^2}$$

is indeed a barrier of width $2n-1$. Positivity is clear, since $x_0 < \frac{1}{4}$ and $1 \leq x \leq 2n-1$. Concavity we obtain by taking the second derivative of u with respect to x: the quantity

$$u''(x) = \frac{-4n^2}{x\sqrt{4n^2-x^2}} \left(\frac{2x_0}{x^2} + \frac{x-x_0}{4n^2-x^2} \right)$$

is always negative within the allowed ranges of the variables. This leaves the two technical assumptions. The first inequality

$$u(1) \geq u(2n-1)$$

is equivalent to

$$(1-x_0)\sqrt{4n^2-1} \geq \frac{2n-1-x_0}{2n-1}\sqrt{4n^2-4n^2+4n-1}$$

or

$$\frac{1-x_0}{2n-1-x_0} \geq \frac{\sqrt{4n-1}}{(2n-1)\sqrt{4n^2-1}}.$$

Since $2n-1 \geq 1$, the quantity on the left is at a minimum when x_0 is as large as possible; we may therefore without loss of generality replace x_0 with the upper bound $\frac{1}{4}$, and attempt to establish

$$\frac{\frac{3}{4}}{2n-\frac{5}{4}} \geq \frac{\sqrt{4n-1}}{(2n-1)\sqrt{4n^2-1}}$$

or

$$\frac{3}{4}(2n-1)\sqrt{4n^2-1} \geq \left(2n-\frac{5}{4}\right)\sqrt{4n-1}.$$

Both sides are positive; squaring them, then, we have to show

$$\frac{9}{16}(2n-1)^2(4n^2-1) \geq \left(2n-\frac{5}{4}\right)^2(4n-1).$$

Multiplying out and pulling all terms to the left, the desired inequality

$$9n^4 - 25n^3 + 24n^2 - 9n + 1 \geq 0,$$

which we rewrite as

$$(n-1)^2(9(n-1)^2 + 11(n-1) + 3) \geq 0,$$

is clearly true for all integers $n \geq 1$. We need also to show that u satisfies Property 3.2. When $n = 1$ (and hence $w = 1$) this is trivial. Otherwise, by the concavity of u, this is equivalent to the inequality

$$u'(1) \leq \frac{1}{2}u(1).$$

The quantity $u(1) - 2u'(1)$ can be written as

$$\frac{1}{\sqrt{4n^2-1}}\left[4n^2(1-3x_0)+1+x_0\right]$$

which is always positive for $n > 1$ and $0 < x_0 < \frac{1}{4}$. Thus u is indeed a barrier, which completes the proof of the theorem, pending verification of some earlier claims. $\square$

3.4 Chambers of any shape

The chambers we have constructed in the proof of Theorem 3.3.1 are, up to an affine transformation, approximately circular—indeed, the formula for u was derived by making the pseudofissures tangent to a circle. It is natural to ask what other contours besides that of a circle may be approximated (to arbitrary precision, and up to an affine transformation) by chambers of lattice polygons. One obvious constraint is convexity. Origin symmetry is another constraint inherent to any technique that depends on the Fissure Slopes Lemma: The maximum slopes through the points $(x_0, y_0 + \frac{1}{2jq})$ and $(x_0, y_0 - \frac{1}{2jq})$ will both approximate $\frac{u(j)}{j-x_0}$, and the minimum slopes through the same pair of points will both approximate $\frac{-\ell(j)}{j-x_0}$. Our next result is that there are no other constraints on the shapes which can be approximated by lattice polygon chambers.

Theorem 3.4.1. *Let X be any convex, compact region of $\mathbb{R}^2$ that satisfies $X = -X$ and such that X has non-empty interior, and choose $0 < \epsilon < 1$. Then there exists a lattice polygon P, a chamber $C \in \mathscr{C}(P)$, and an affine transformation L of $\mathbb{R}^2$ such that $(1-\epsilon)L(C) \subset X \subset (1+\epsilon)L(C)$.*

It will be convenient to normalize X to lie within the square $-1 \leq x, y \leq 1$ and to further assume that it includes the points $(\pm 1, 0)$ and $(0, \pm 1)$. This is justified by the following proposition and its corollary, which guarantees an affine image of X that fits such a normalization.

We begin with a general proposition (probably classical) about compact convex regions in any dimension.

Proposition 3.4.2. Let $X \subset \mathbb{R}^d$ be a compact, convex region with non-empty interior. Then there exists an affine transformation A of $\mathbb{R}^d$ such that $A(X)$ lies within the closed unit cube $0 \leq x_1, x_2, \ldots, x_d \leq 1$ but such that, for each of the standard basis vectors $e_i \in \{e_1, \ldots, e_d\}$, $A(X)$ has non-empty intersection with the translate $A(X) + e_i$.

Proof. Consider the set $\mathcal{A}$ of all affine transformations for which the image of X lies within the closed unit cube. Since X is compact with non-empty interior, and the closed unit cube is compact, the set $\mathcal{A}$ is also compact. The determinant of the linear portion of an affine transformation is a continuous function of the transformation chosen; let $A \in \mathcal{A}$ be a transformation for which this determinant is maximized over $\mathcal{A}$. It follows that the volume of the image of X under any transformation in $\mathcal{A}$ cannot exceed the volume of $A(X)$. Any affine transformation $C \in \mathcal{A}$ for which the volume of $C(X)$ is maximal must cause X to touch each of the facets of the unit cube, for if $C(X)$ were disjoint from some hyperplane $x_i = 0$ or $x_i = 1$, then $C(X)$ could be composed with an affine stretch in the corresponding direction, keeping the image within the unit cube, but yielding a greater volume.

Now suppose that $A(X)$ is disjoint from $A(X) + e_i$ for some standard basis vector e_i. Since $A(X)$ and $A(X) + e_i$ are convex, compact, and disjoint, there is some separating hyperplane h which is disjoint from either of them. Let B be the affine transformation which fixes each coördinate x_j when $j \neq i$ but which maps h to the hyperplane $x_i = 1$ and maps $h - e_i$ to the hyperplane $x_i = 0$. The image of the unit cube under the inverse of B is a parallelotope of volume 1, and so B preserves volume. In particular, $BA(X)$ has the same volume as $A(X)$ and lies within the unit cube, but $BA(X)$ is disjoint from the facets $x_i = 0$ and $x_i = 1$. This is a contradiction, and we conclude that $A(X)$ must intersect each of the translates $A(X) + e_i$. □

Specializing to two dimensions and central symmetry, we obtain the desired corollary.

Corollary 3.4.3. Without loss of generality, the region X in the statement of Theorem 3.4.1 can be assumed to lie within the square $-1 \leq x, y \leq 1$ and to contain the points $(\pm 1, 0)$ and $(0, \pm 1)$.

Proof. We are given $X \subset \mathbb{R}^2$ a compact, convex region with non-empty interior. By Proposition 3.4.2, X is affinely equivalent to a region $A(X)$ lying within the closed square $0 \leq x, y \leq 1$ such that $A(X)$ intersects the two translates either one unit above itself or one unit to the right. In other words, there exist real numbers $0 \leq a, b \leq 1$ such that the four points

$$(0, a),\ (1, a),\ (b, 0),\ \text{and } (b, 1)$$

all lie within $A(X)$. The region X in the statement of Theorem 3.4.1 satisfies $X = -X$, and as a result the affine image $A(X)$ must have symmetry about some point $A(\mathbf{0}) =$

$(\frac{h}{2}, \frac{k}{2})$ and hence must contain the four points

$$(h-0, k-a),\ (h-1, k-a),\ (h-b, k-0),\ \text{and}\ (h-b, k-1).$$

Since $A(X)$ lies within the unit square, we must have $h = 1$ and $k = 1$. In the generic case the points $(0, a)$ and $(1-1, 1-a)$ will be identical, implying $a = \frac{1}{2}$, but in any case by convexity the midpoint $(0, \frac{1}{2})$ of these two points will be contained in $A(X)$. The same argument applied to each of the four sides of the unit square allows us to assume without loss of generality that $a = \frac{1}{2}$ and $b = \frac{1}{2}$. If D is the affine transformation which doubles the size of $A(X)$ and repositions it within the square $-1 \leq x, y \leq 1$, then $DA(X)$ will contain the points $(\pm 1, 0)$ and $(0, \pm 1)$. By absorbing the affine transformation DA into the affine transformation L in the statement of Theorem 3.4.1, we may assume without loss of generality that X itself lies within the square $-1 \leq x, y \leq 1$ and contains the points $(\pm 1, 0)$ and $(0, \pm 1)$. □

Proof of Theorem 3.4.1. Let X be any convex, compact region of $\mathbb{R}^2$ that satisfies $X = -X$ and such that X has non-empty interior. Since we are only concerned with the shape of X up to affine transformation, Corollary 3.4.3 allows us to assume without loss of generality that X lies within the square $-1 \leq x, y \leq 1$ and contains the points $(\pm 1, 0)$ and $(0, \pm 1)$. By passing to an approximation (and adjusting the value of ϵ accordingly) we make the further simplifying assumption that ∂X is a smooth curve with everywhere positive curvature.

Consider the second quadrant of ∂X as the graph of a function $b : [-1, 0] \to [0, 1]$ with the following properties:

- b is continuous on the interval $[-1, 0]$.
- $b(-1) = 0$
- $b(0) = 1$
- $b'(t) > 0$ for $-1 < t < 0$
- $\lim_{t \to -1^+} b'(t) = \infty$
- $b'(0) = 0$
- $b''(t) < 0$ for $-1 < t \leq 0$.

Given a point (x_0, y_0) and constants n and q, we let $\widehat{X}$ be the the region of the same shape as X, but of $\frac{1}{4nq}$ the size, and centered at the point (x_0, y_0). We will abuse notation by referring to the image of the second quadrant of the boundary of X (for example) as the "second quadrant of $\partial\widehat{X}$," ignoring the fact that $\widehat{X}$ will typically lie entirely within the first quadrant, or the fourth. We also note, by way of apology, that the notation $\widehat{X}$ suppresses the dependence on x_0, y_0, n, and q.

Just as $b(t)$ traces the second quadrant of X, we will denote by $\widehat{b}(x)$ that function whose graph is the second quadrant of $\partial\widehat{X}$. To be precise, we define $\widehat{b} : [x_0 - \frac{1}{4nq}, x_0] \to [y_0, y_0 + \frac{1}{4nq}]$ in such a way that

$$\widehat{b}\left(x_0 + \frac{t}{4nq}\right) = y_0 + \frac{b(t)}{4nq}.$$

Our goal is use the function $b(t)$ to construct, on some sufficiently large interval $[1, 2n-1]$, a barrier u whose j^{th} pseudofissures through the series of points

$$(x_0, y_j), \quad y_j = y_0 + \frac{1}{2jq}$$

are each tangent to the graph $y = \widehat{b}(x)$. By symmetry, this will also ensure that the j^{th} pseudofissures through the points

$$(x_0, \bar{y}_j), \quad \bar{y}_j = y_0 - \frac{1}{2jq}$$

are similarly tangent to the fourth quadrant of $\partial\widehat{X}$. The same process will give us a lower barrier ℓ if we apply the construction, not to $b(t)$, but rather to the function whose negative defines the third quadrant of ∂X, and the j^{th} pseudofissures of $-\ell$ through the points (x_0, y_n) and $(x_0, \bar{y}_n)$ will approximate the third and first quadrants, respectively, of $\partial\widehat{X}$. Since the procedures for constructing u and ℓ are identical except for the choice of input function, we will consider only the input function $b(t)$ and its consequent barrier $u(x)$, without explicitly mentioning $\ell(x)$ except as necessary to establish the condition $\frac{5}{6} \leq \frac{u(1)}{\ell(1)} \leq \frac{6}{5}$. The construction of both an upper and a lower barrier is, however, implied, and in particular any constraints which require n to be "sufficiently large" are to be understood as constraints coming simultaneously from the construction of u and of ℓ; whichever requirement is stricter must be the one applied.

The width $2n-1$ of the upper and lower barriers u and ℓ is an arbitrary parameter during the construction of the barriers themselves, but it must eventually be chosen to be large enough that the resulting pseudochamber D approximates $\widehat{X}$ (or rather approximates, to be pedantic, the smooth approximation of $\widehat{X}$ that was chosen at the beginning of the proof, claiming no loss of generality) to the desired degree of accuracy. The Pseudochamber Approximation Lemma then guarantees the existence of some lattice polygon P and a chamber $C \in \mathscr{C}(P)$ that approximates, up to affine transformation and again to within arbitrary precision, this approximation D of an approximation of $\widehat{X}$. Assuming that tolerances are chosen appropriately at each stage of approximation, a lattice polygon chamber can therefore be produced to satisfy the statement of the theorem for any specified value of ϵ.

Although the pseudofissures of u form a discrete set of lines indexed by integers $1 \leq j \leq 2n-1$ (odd integers, in the case of interest), we will extend these pseudofissures to a continuous family $\mathcal{L}$ of lines of slope $\frac{u(x)}{x-x_0}$. Whatever the value of n is, it is important to know where x lies relative to the range $[0, 2n]$, and we introduce the continuous parameter $s = \frac{x}{2n}$ as a measure of this. Instead of passing through one of the discrete set of points (x_0, y_j), a line $L \in \mathcal{L}$ of slope $\frac{u(2ns)}{2ns-x_0}$ must pass through the point

$$\left(x_0, y_0 + \frac{1}{4nq}\frac{1}{s}\right).$$

We wish to construct $u(x)$ in such a manner that each line $L \in \mathcal{L}$ is also tangent to the graph of the function $\widehat{b}(x)$.

The construction of the pseudochamber happens on a very small scale, $\frac{1}{4nq}$ the size of X, whereas the barrier $u(j)$ is as large as the pseudochamber is small; at both scales, however, the slope is the same. Instead of defining $u(x)$ directly in terms of $\widehat{b}(x)$, it is convenient to relate the functions $b(t)$ and $u(x)$ to each other by means of an intermediate function $m(s)$ which records, for each value of the parameter s, the slope of the corresponding line in $\mathcal{L}$. The value of $m(s)$ can be defined (on an intermediate scale) with direct reference to the function $b(t)$, and the three functions are related as follows:

$$b'(t) = m(s) = \frac{u(x)}{x - x_0}.$$

To complete the definition of $u(x)$ based on $b(t)$, we must relate not only the function values but also the parameters. We have, by definition, $s = \frac{x}{2n}$, but the relationship between the parameters s and t is slightly more subtle, as illustrated by Figure 3.3: The point $\left(x_0, y_0 + \frac{1}{4nq}\frac{1}{s}\right)$ corresponds to the point $\left(0, \frac{1}{s}\right)$ in the graph of $b(t)$, which by tangency constraints must be the same as the point $(0, b(t) - tb'(t))$. We thus have

$$b(t) - tb'(t) = \frac{1}{s} = \frac{2n}{x}.$$

Figure 3.3: The function $b(t)$, and how it determines $m(s)$

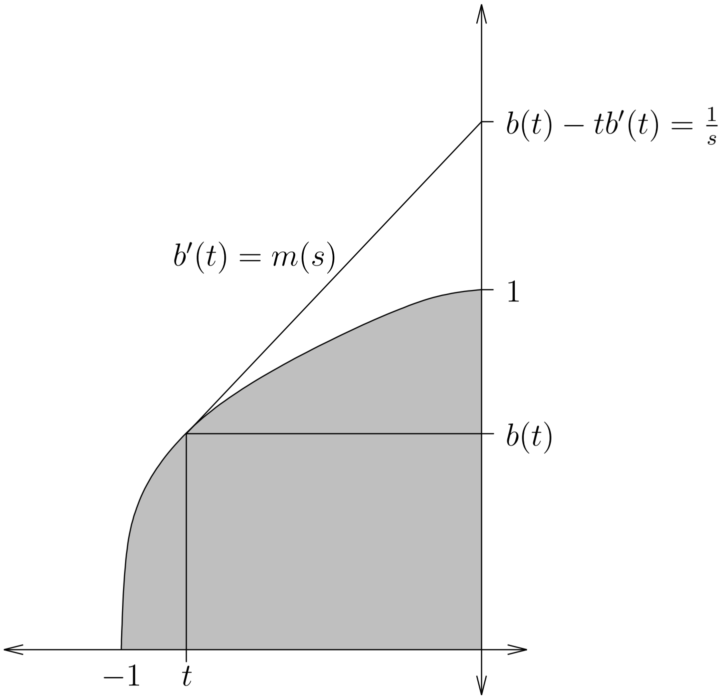

Note that each of the three parameters $t \in [-1, 0]$, $s \in \left[\frac{1}{2n}, \frac{2n-1}{2n}\right]$, and $x \in [1, 2n-1]$ is monotone increasing with respect to the other two parameters.

This gives us well-defined functions u and ℓ, but before we can apply the (still to be proven) Pseudochamber Approximation Lemma, we must show that u and ℓ satisfy all the requirements of barriers—which involves, as in the proof of Theorem 3.3.1, a series of straightforward but cumbersome calculations—and also show that $\frac{5}{6} \leq \frac{u(1)}{\ell(1)} \leq \frac{6}{5}$.

The last detail we take care of first. Referring again to Figure 3.3, it is clear that the tangent line, of slope $m(s)$, must intersect the vertical line $t = -1$ at a height somewhere between 0 and 1, and it intersects the line $t = 0$ at height exactly $\frac{1}{s}$. This means that the inequality

$$\frac{1}{s} - 1 \leq m(s) \leq \frac{1}{s}$$

must hold for any value of s. Applying this to $m\left(\frac{1}{2n}\right) = \frac{u(1)}{1-x_0}$ gives us

$$2n - 1 \leq \frac{u(1)}{1 - x_0} \leq 2n.$$

This inequality is independent of $b(t)$; in particular, the value of $\ell(1)$ is constrained to lie within the same narrow range, depending only on n and x_0, in which $u(1)$ is found.

It follows that we must have $\frac{5}{6} \le \frac{u(1)}{\ell(1)} \le \frac{6}{5}$ whenever $n \ge 3$.

Now to show that u (and similarly that ℓ) is indeed a barrier. Positivity is straightforward; since $\frac{1}{s} > 1$, we have

$$0 < \frac{1}{s} - 1 \le m(s) = \frac{u(x)}{x - x_0}.$$

Postponing concavity for the time being, we establish the first of the two technical assumptions. Property 3.1 states

$$u(2n-1) \le u(1).$$

This may not be true if n is not sufficiently large. In particular, we need to ensure that the slope $m\big(\frac{2n-1}{2n}\big)$ of the last pseudofissure is sufficiently small—it needs to have slope no larger than $\frac{3}{4}$. The fact that $b'(t)$ decreases in a strictly monotone fashion from infinity to zero allows us to specify a (unique) number $t_0 < 0$ such that $b'(t_0) = \frac{3}{4}$, and we let s_0 be the corresponding value of s. In other words, referring once more to Figure 3.3, we set $\frac{1}{s_0} = b(t_0) - t_0 b'(t_0)$, with the assurance that s_0 is then strictly less than 1. As long as n is chosen to be sufficiently large that $\frac{2n-1}{2n} \ge s_0$, we will then be assured that $m\left(\frac{2n-1}{2n}\right) \le m(s_0) = \frac{3}{4}$. (The same step in the construction of the lower barrier ℓ will give a separate and possibly stronger constraint on how large n must be.) With this assumption in hand, we have

$$\frac{u(2n-1)}{2n-1-x_0} = m\bigg(\frac{2n-1}{2n}\bigg) \le \frac{3}{4}$$

which together with the already established inequality

$$2n - 1 \le \frac{u(1)}{1 - x_0}$$

gives us

$$u(2n-1) \le \frac{3}{4}(2n-1-x_0) < \frac{3}{4}(2n-1) < (1-x_0)(2n-1) \le u(1).$$

It is a strange and useful fact that the value of $m'(s)$ depends only on the parameters s and t, and not explicitly on the function $b(t)$. On the one hand, starting with the definition $m(s) = b'(t)$ and differentiating implicitly, we have

$$m'(s)\, ds = b''(t)\, dt.$$

But if we start instead with the equation $\frac{1}{s} = b(t) - tb'(t)$, implicit differentiation then gives us

$$\frac{-1}{s^2}\, ds = b'(t)\, dt - dt\, b'(t) - tb''(t)\, dt.$$

Combining these, we obtain

$$m'(s)\,ds = b''(t)\,dt = \frac{1}{ts^2}\,ds$$

or in other words

$$m'(s) = \frac{1}{ts^2}.$$

Since t is always negative, this means, for example, that $m(s)$ always decreases as s increases (or as $\frac{1}{s}$ decreases), which is true regardless of the fact that $b''(t) < 0$.

We now establish Property 3.2, in the guise (under the still-postponed assumption of concavity) of the inequality

$$\frac{u'(1)}{u(1)} \leq \frac{1}{2}.$$

We start with the definition

$$u(x) = (x - x_0)m(s).$$

Differentiation with respect to x gives us

$$\begin{aligned} u'(x) &= m(s) + (x - x_0)m'(s)\frac{ds}{dx} \\ &= m(s) + \frac{x - x_0}{ts^2}\frac{1}{2n} \\ &= m(s) + \frac{x - x_0}{tsx} \end{aligned}$$

and dividing by $u(x)$ we obtain

$$\frac{u'(x)}{u(x)} = \frac{1}{x - x_0} + \frac{1}{m(s)tsx}.$$

Since $x - x_0$ is at least $\frac{3}{4}$ and since $m(s)$ is at most $\frac{1}{s}$ this gives us the inequalities

$$\frac{u'(x)}{u(x)} \leq \frac{4}{3} + \frac{1}{tx} < \frac{4}{3} - \frac{1}{x}$$

(since t lies strictly between -1 and 0) and in particular

$$\frac{u'(1)}{u(1)} \leq \frac{4}{3} - \frac{1}{1} < \frac{1}{2}$$

as was to be shown.

This leaves only the question of the concavity of $u(x)$ in order to show that u (and likewise that ℓ) is a barrier to which we can apply the Pseudochamber Approximation Lemma. We will show that $u''(x) < 0$ holds within the allowed range. We have already computed

$$u'(x) = m(s) + \frac{x - x_0}{tsx} = m(s) + \frac{1}{ts} - \frac{x_0}{tsx}.$$

Differentating once more with respect to x we obtain

$$u''(x) = m'(s)\frac{ds}{dx} - \frac{1}{t^2 s}\frac{dt}{dx} - \frac{1}{ts^2}\frac{ds}{dx} + \frac{x_0}{t^2 sx}\frac{dt}{dx} + \frac{x_0}{ts^2 x}\frac{ds}{dx} + \frac{x_0}{tsx^2}.$$

Substituting for $m'(s)$ once more and regrouping, we have

$$\begin{aligned} u''(x) &= -\frac{x-x_0}{t^2 sx}\frac{dt}{dx} + \frac{x_0}{ts^2 x}\frac{ds}{dx} + \frac{x_0}{tsx^2} \\ &= -\frac{x-x_0}{t^2 x^2}\frac{dt}{ds} + \frac{2x_0}{tsx^2}. \end{aligned}$$

Both summands are negative. It follows that u is concave downward, and thus that u is a barrier.

Remark 3.4.4. In some sense we have been working backwards here; any valid barrier $u(x)$ determines some function $b(t)$, and hence the shape of a pseudochamber, and, in the limit as n and N go to infinity, a region X. The function $b(t)$ so derived will not necessarily satisfy all of originally assumed conditions, however, because for $u(x)$ to be a barrier is a strictly weaker condition. Let us examine in particular what happens when $u(x)$ is only weakly concave, or in other words when $u''(x) = 0$. According to the last equation we have derived, this forces $\frac{ds}{dt}$ to be negative, or in other words it forces X to be strictly non-convex. When this happens, the lines of tangency overlap in the interior of the non-convex region $\widehat{X}$, and in fact for $u(x)$ a constant function (in other words, for P a rectangle), every such pseudochamber will in fact be an octagon, no matter how large n is. In this case the actual chambers centered at points (x_0, y_0) of the class that we have studied will eventually converge to octagons as well, as N increases, and any such chamber with more than 8 edges is in some sense an "artifact" of the slope errors of size at most $q + \frac{4}{3}|y_j|$. It is not entirely surprising, in light of this analysis, that no chamber of a rectangle has yet been found with more than 15 edges. While it is quite possible that some different class of points yields chambers of arbitrary complexity even within a rectangle, it remains at least possible that there is a uniform bound on the complexity of a chamber complex in terms of the number of edges of the lattice polygon itself.

Having shown that u and ℓ are barriers, and that $\frac{5}{6} \leq \frac{u(1)}{\ell(1)} \leq \frac{6}{5}$, we can apply the Pseudochamber Approximation Lemma to the lattice polygon with upper barrier Nu and lower barrier $N\ell$, for large N, to find a chamber arbitrarily close to an affine image of the pseudochamber whose $4n$ pseudofissures are each tangent to the boundary of $\widehat{X}$, or more precisely to the smooth approximation of $\widehat{X}$ chosen at the first stage of the proof.

The parameter N thus ensures sufficient approximation to the specified pseudochamber, and the pseudochamber itself can be made to approximate $\widehat{X}$ to the desired degree of accuracy by a sufficiently high value of n. Since the accuracy obtained at any one stage is arbitrary, the statement of the theorem is satisfied for any specified value of ϵ. $\square$

3.5 Nine more pages of inequalities

We have to this point postponed the details of the proofs of the Fissure Slopes Lemma and the Pseudochamber Approximation Lemma, which situation we now remedy.

3.5.1 Fissure slopes in detail

Proof of the Fissure Slopes Lemma. For any positive integer w, choose a positive integer β and an odd positive integer q such that $w < 2^\beta < \frac{1}{4}q$, and set $x_0 = \frac{2^\beta}{q}$. Choose j an odd integer with $1 \leq j \leq w$, and let y_j be a multiple of $\frac{1}{jq}$ such that y_j is not a multiple of $\frac{1}{iq}$ for any odd positive integer $i < j$. Let u and ℓ be barriers of width w satisfying $u(1), \ell(1) \geq 16|y_j|$; $u(w), \ell(w) \geq wq - 2^\beta$; and $\frac{5}{6} \leq \frac{u(1)}{\ell(1)} \leq \frac{6}{5}$. Let P be the polygon with upper barrier u and lower barrier ℓ. For each integer k in the range $1 \leq k \leq w$, denote by M_k the set of all slopes of fissures passing through the point (x_0, y_j) whose right endpoint has k as its x-coördinate, and define

$$M = \bigcup_{k=1}^{w} M_k.$$

To prove Lemma 3.2.2, we must demonstrate that

1. M is non-empty: some fissure of P passes through (x_0, y_j),
2. The maximum value of M is within $q + \frac{4}{3}|y_j|$ of $\frac{u(j)}{j-x_0}$, and
3. The minimum value of M is within $q + \frac{4}{3}|y_j|$ of $\frac{-\ell(j)}{j-x_0}$.

Rather than dealing with M as a whole, however, we will consider each of the sets M_k separately. Accordingly, for each integer k in the range $1 \leq k \leq w$, factor k as a power of two times an odd integer

$$k = 2^\alpha i.$$

The values of α and i are unique given k, and α is always strictly less than β, as before. When i is less than j (and in particular when $k < j$) we know by assumption that y_j is not a multiple of $\frac{1}{iq}$, and Lemma 3.2.1 ensures that M_k is empty. If on the other hand $j \leq i$ (implying $j \leq k$) then we have, as a straightforward consequence of inequality 3.4 for barriers,

$$\frac{u(j)}{j - x_0} \geq \frac{u(k)}{k - x_0}$$

and

$$\frac{-\ell(j)}{j - x_0} \leq \frac{-\ell(k)}{k - x_0}.$$

It thus suffices to prove that whenever k takes a value such that y_j is a multiple of $\frac{1}{iq}$ (and in particular when $k = j = i$),

1′. M_k is non-empty,

2′. The maximum value of M_k is within $q + \frac{4}{3}|y_j|$ of $\frac{u(k)}{k-x_0}$, and

3′. The minimum value of M_k is within $q + \frac{4}{3}|y_j|$ of $\frac{-\ell(k)}{k-x_0}$,

since the range of slopes for k lies within the range of slopes for j when $j \leq k$, and the desired error bound $q + \frac{4}{3}|y_j|$ is independent of k.

Suppose then that $y_j = \frac{c}{iq}$, where c is an integer. Referring to Equation 3.6 from the proof of Lemma 3.2.1, we see that if a and b are integers in the ranges $\frac{-\ell(1)}{2} \leq a \leq \frac{u(1)}{2}$ and $-\ell(k) \leq b \leq u(k)$ that satisfy the equation

$$(b - a)2^{\beta-\alpha} + aiq = c,$$

which we rewrite as

$$a\left(iq - 2^{\beta-\alpha}\right) + b\left(2^{\beta-\alpha}\right) = c, \tag{3.7}$$

then the fissure of P with endpoints $(0, a)$ and (k, b) passes through the point (x_0, y_j), and M_k contains the slope $\frac{b-a}{k}$. The implication goes both ways; every value of M_k comes from such a pair (a, b). The odd number $iq - 2^{\beta-\alpha}$ is relatively prime to $2^{\beta-\alpha}$, so let s and t be a pair of integers satisfying

$$s(iq - 2^{\beta-\alpha}) + t(2^{\beta-\alpha}) = 1$$

and define, for every integer n,

$$\begin{aligned} a_n &= cs - n(2^{\beta-\alpha}) \\ b_n &= ct + n(iq - 2^{\beta-\alpha}). \end{aligned}$$

The pairs (a_n, b_n) are exactly the solutions to Equation 3.7, and thus M_k consists of those rational numbers

$$\frac{b_n - a_n}{k} = \frac{ct - cs + niq}{i2^\alpha} = \frac{t-s}{k}c + n\frac{q}{2^\alpha}$$

for which $(0, a_n)$ and (k, b_n) are lattice points of P. Since a_n and b_n each depend monotonically on n, the values of n that satisfy both of these constraints form a set of consecutive integers. It follows that the slopes M_k form a finite arithmetic progression of rational numbers spaced $\frac{q}{2^\alpha}$ apart.

The remainder of the proof of Lemma 3.2.2 is concerned with a series of inequalities, the first few of which are useful in the pursuit of the claim that $\frac{b_n - a_n}{k}$ belongs to M_k if and only if $-\ell(k) \leq b_n \leq u(k)$, or in other words

$$-\ell(k) \leq b_n \leq u(k) \;\Rightarrow\; -\frac{1}{2}\ell(1) \leq a_n \leq \frac{1}{2}u(1)$$

regardless of k. Rather than dealing with both barriers simultaneously, we prove instead the pair of implications

$$|b_n| \leq u(k) \;\Rightarrow\; |a_n| \leq \frac{5}{12}u(1)$$

and

$$|b_n| \leq \ell(k) \;\Rightarrow\; |a_n| \leq \frac{5}{12}\ell(1)$$

which together establish the desired claim once we recall the assumption $\frac{5}{6} \leq \frac{u(1)}{\ell(1)} \leq \frac{6}{5}$. In fact we give only the proof of the implication involving the upper barrier, since the assumptions on u and ℓ are completely symmetric.

First we need to establish a few inequalities that depend only on the bounds $0 < x_0 < \frac{1}{4}$ and $k \geq 1$. Since the quantity

$$\frac{k}{k - x_0}$$

is monotone increasing with respect to x_0 and monotone decreasing with respect to k, it approaches its maximum as x_0 approaches $\frac{1}{4}$ and $k = 1$, and

$$\frac{k}{k - x_0} < \frac{4}{3}. \tag{3.8}$$

For the same reasons,

$$\frac{k+1}{k - x_0} < \frac{8}{3}$$

and thus

$$\frac{x_0}{k - x_0}(k+1) < \frac{2}{3}. \tag{3.9}$$

Now, regardless of whether or not $(0, a_n)$ and (k, b_n) are lattice points of P, this pair of points will always be collinear with the point (x_0, y_j), and thus

$$\frac{b_n - y_j}{k - x_0} = \frac{y_j - a_n}{x_0},$$

which reduces to

$$a_n = \frac{-x_0}{k - x_0} b_n + \frac{k}{k - x_0} y_j,$$

and hence

$$|a_n| \leq \frac{x_0}{k - x_0} |b_n| + \frac{k}{k - x_0} |y_j|.$$

Assuming, then, that $|b_n| \leq u(k)$, and finally invoking the mysterious assumption that $|y_j| \leq \frac{u(1)}{16}$, we have

$$|a_n| \leq \frac{x_0}{k - x_0} u(k) + \frac{k}{k - x_0} \frac{u(1)}{16}.$$

Applying the technical condition (3.2) for barriers, we have

$$|a_n| \leq \frac{x_0}{k - x_0} (k+1) \frac{u(1)}{2} + \frac{k}{k - x_0} \frac{u(1)}{16}$$

to which we can apply the inequalities (3.8) and (3.9) to obtain

$$|a_n| \leq \frac{2}{3} \frac{u(1)}{2} + \frac{4}{3} \frac{u(1)}{16}$$

and finally

$$|a_n| \leq \frac{5}{12} u(1).$$

We have shown that M_k consists of rational numbers spaced $\frac{q}{2^\alpha}$ apart, with only the constraint that $-\ell(k) \leq b_n \leq u(k)$. Up to an error of size less than $\frac{q}{2^\alpha}$, then, and hence of size less than q, the slopes $m \in M_k$ are exactly in the range

$$\frac{-\ell(k) - y_j}{k - x_0} \leq m \leq \frac{u(k) - y_j}{k - x_0}$$

and it follows that M_k is non-empty so long as $u(k) + \ell(k) \geq q(k - x_0) = kq - 2^\beta$. By assumption we have $u(w) + \ell(w) \geq 2(wq - 2^\beta)$, and hence $u(w) + \ell(w) \geq kq - 2^\beta$, but $u(w)$ is a lower bound for all $u(k)$ (since $u(1) \geq u(w)$ and u is concave downward) and similarly $\ell(k) \geq \ell(w)$ regardless of k. It follows that M_k is non-empty, and we have proven Statement $1'$ for all values of k such that y_j is a multiple of $\frac{1}{iq}$, as was desired. Since $k - x_0 > \frac{3}{4}$, the continuous bounds on slope are within $\frac{4}{3}|y_j|$ of $\frac{u(k)}{k - x_0}$ and $\frac{-\ell(k)}{k - x_0}$ respectively, and the actual greatest and least values of M_k are within q (or within $\frac{q}{2^\alpha}$, more precisely) of the continuous bounds, which completes the proof of Statements 2′ and 3′ under the same assumption that y_j is a multiple of $\frac{1}{iq}$. Since M_k is either empty or bounded strictly between $\frac{-\ell(j)}{j-x_0} - q - \frac{4}{3}|y_j|$ and $\frac{u(j)}{j-x_0} + q + \frac{4}{3}|y_j|$ for $k < j$ or $k > j$, we may apply these statements to the special case $k = j$ to obtain Statements 1, 2, and 3. This completes the proof of Lemma 3.2.2. □

3.5.2 Pseudochamber approximation in detail

Proof of the Pseudochamber Approximation Lemma. Let w be any positive integer, and choose β and q such that $w < 2^\beta < \frac{1}{4}q$. As in the statement of the Pseudochamber Approximation Lemma (or equivalently, as in the statement of the Fissure Slopes Lemma) let $x_0 = \frac{2^n}{q}$. Let y_0 be any odd multiple of $\frac{1}{2q}$, which we may write as $\frac{c}{q} + \frac{1}{2q}$ for some integer c. For each odd integer in the range $1 \leq j \leq w$, the two points

$$y_j = y_0 + \frac{1}{2jq} = \frac{c}{q} + \frac{j+1}{2jq} \quad \text{and} \quad \overline{y}_j = y_0 - \frac{1}{2jq} = \frac{c}{q} + \frac{j-1}{2jq}$$

are the closest multiples of $\frac{1}{jq}$ to y_0, and furthermore are not multiples of $\frac{1}{iq}$ for any odd positive integer $i < j$, a required assumption for the Fissure Slopes Lemma. The maximum absolute value of y_j or $\overline{y}_j$ for any j is of size $\frac{|c|+1}{q}$, so the error bound from the Fissure Slopes Lemma will be of size no more than

$$E = q + \frac{4}{3}\frac{|c|+1}{q},$$

and we note that E depends only on x_0 and y_0. Let u and ℓ be barriers such that $\frac{5}{6} \leq \frac{u(1)}{\ell(1)} \leq \frac{6}{5}$, and let N be a positive integer sufficiently large to ensure that

$$Nu(1) = N\ell(1) \geq 16\frac{|c|+1}{q}$$

and to ensure both that $Nu(w) > 2wE$ and that $N\ell(w) > 2wE$, which in turn easily imply

$$Nu(w) > wE > wq > wq - 2^\beta$$

and similarly for $\ell(w)$. Let $P(N)$ represent the polygon with upper barrier Nu and lower barrier $N\ell$, and apply Lemma 3.2.2 to obtain information about the slopes of fissures of $P(N)$ at each of the points (x_0, y_j) and $(x_0, \overline{y}_j)$. We will demonstrate that this information is sufficient to completely specify the shape of the chamber $C(N)$ of $P(N)$ containing the point (x_0, y_0). (If some fissure contained the point (x_0, y_0), this chamber would not be uniquely determined, but that cannot happen, by Lemma 3.2.1, because $y_0 = \frac{2c+1}{2q}$ is not a multiple of $\frac{1}{iq}$ for any odd integer i.) For each value of N and of j, let $M(N,j)$ and $m(N,j)$ represent the maximum and minimum slopes, respectively, of fissures of $P(N)$ through the point (x_0, y_j), and similarly let $\overline{M}(N,j)$ and $\overline{m}(N,j)$ represent the maximum and minimum slopes through $(x_0, \overline{y}_j)$. Denote by $Q(N,j)$ the quadrilateral bounded by these four fissures and containing the point (x_0, y_0). By Lemma 3.2.2, $M(N,j)$ and $\overline{M}(N,j)$ are both within E of $\frac{Nu(j)}{j-x_0}$, and $m(N,j)$ and $\overline{m}(N,j)$ are both within E of $\frac{-N\ell(j)}{j-x_0}$. We have chosen N sufficiently large that

$$Nu(j) \geq Nu(w) > 2wE > 2(j-x_0)E,$$

which gives us

$$M(N,j), \overline{M}(N,j) \geq \frac{Nu(j)}{j-x_0} - E > E \tag{3.10}$$

and the same argument gives us

$$m(N,j), \overline{m}(N,j) \leq \frac{-N\ell(j)}{j-x_0} + E < -E \tag{3.11}$$

which in particular tells us that the maximum and minimum slopes are positive and negative, respectively, at either point. From the quadrilaterals $Q(N,j)$ we construct another set of quadrilaterals that have been stretched by a factor of N in the x-direction: let $\widehat{Q}(N,j)$ be the image of $Q(N,j)$ under the transformation

$$(x,y) \mapsto (x_0 + N(x-x_0), y).$$

Each region $\widehat{Q}(N,j)$ is also bounded by line segments that pass through the points (x_0, y_j) and $(x_0, \overline{y}_j)$, but the slopes after rescaling are within $\frac{E}{N}$ of $\frac{u(j)}{j-x_0}$ or within $\frac{E}{N}$ of $\frac{-\ell(j)}{j-x_0}$. If we define R_j to be the parallelogram bounded by lines through (x_0, y_j) and $(x_0, \overline{y}_j)$ of slopes exactly $\frac{u(j)}{j-x_0}$ and $\frac{-\ell(j)}{j-x_0}$, then we have, for each odd $1 \leq j \leq w$,

$$\lim_{N\to\infty} \widehat{Q}(N,j) = R_j.$$

What is more, we have

$$\bigcap_j R_j = D,$$

where D is the pseudochamber in the statement of the theorem. Combining these results, we have

$$\lim_{N\to\infty} \bigcap_j \widehat{Q}(N,j) = D.$$

To prove the theorem, then, it remains only to show that, for N sufficiently large,

$$C(N) = \bigcap_j Q(N, j),$$

which would imply

$$\lim_{N\to\infty} \widehat{C}(N) = D,$$

where $\widehat{C}(N)$, as in the statement of the theorem, is the image of $C(N)$ under the image of the same affine transformation

$$(x, y) \mapsto (x_0 + N(x - x_0), y)$$

that sent $Q(N, j)$ to $\widehat{Q}(N, j)$. Clearly we have

$$C(N) \subset \bigcap_j Q(N, j),$$

since each of the quadrilaterals $Q(N, j)$ is bounded by fissures of $P(N)$ and contains the point (x_0, y_0). We need only to show, for N sufficiently large, that no fissure of $P(N)$ passes through the interior of $\bigcap_j Q(N, j)$.

Suppose by way of contradiction that F is a fissure of $P(N)$ that passes through the interior of each of the quadrilaterals $Q(N, j)$. It is useful to have upper and lower bounds on the x-coördinates of points within $Q(N, j)$. Using the bounds given by the inequalities 3.10 and 3.11 on the slopes of the four sides of $Q(N, j)$, together with the fact that $y_j - \overline{y}_j = \frac{1}{jq}$, we have

$$|x - x_0| < \frac{1}{2Ejq}$$

for any point $(x, y) \in Q(N, j)$. In particular, each of the sets $Q(N, j)$ lies entirely within the strip $0 < x < 1$, and so F, which crosses through the interior of each of them, must have one endpoint on the y-axis and another endpoint with x-coördinate k for some integer $k \geq 1$. Such a fissure must pass through exactly one point with x-coördinate equal to x_0; if we call that point (x_0, y_*), then Lemma 3.2.1 tells us that y_* must be a multiple of $\frac{1}{iq}$, where i is the odd part of k. If F passes through the interior of each quadrilateral $Q(N, j)$, it must in particular pass through the interior of $Q(N, i)$, which means that y_* must be strictly greater than or strictly less than y_i. Up to this point we have not made use of the symmetric assumptions on the barriers u and ℓ, so we may, without loss of generality, assume $y_* - y_1 = \frac{d}{iq}$, where d is a positive integer. We need bounds on the slope of F, which we obtain by asking where F touches the vertical line

$x = i$ (as it must, since $i \leq k$). If F contains some point (i, b) that lies within the lattice polygon $P(N)$, we must have $-\ell(i) \leq b \leq u(i)$. Let

$$m_F = \frac{b - y_*}{i - x_0}$$

represent the slope of the fissure F, then we have the following constraints:

$$\frac{-\ell(i) - y_*}{i - x_0} \leq m_F \leq \frac{u(i) - y_*}{i - x_0}.$$

There are two possible cases: either F crosses through $Q(N, i)$ at an x-value less than x_0 (but greater than $x_0 - \frac{1}{2Eiq}$), or F crosses through $Q(N, i)$ at an x-value greater than x_0 (but less than $x_0 - \frac{1}{2Eiq}$). In the first case, F crosses the fissure of slope $M(N, i)$; in the second case, it crosses the fissure of slope $m(N, i)$. In either case, let m_Q represent the slope of the fissure crossed. If m_Q is positive, then m_F is more positive; if m_Q is negative, then m_F is more negative. More precisely, we must have

$$|m_F| - |m_Q| > \frac{d}{iq} / \frac{1}{2Eiq}$$

or

$$|m_F| - |m_Q| > 2Ed.$$

In the first case, where m_F and m_Q are both positive, we replace m_F with its upper bound and m_Q with the lower bound given by Lemma 3.2.2 and obtain

$$\frac{u(i) - y_*}{i - x_0} - \left(\frac{u(i)}{i - x_0} - E \right) > 2Ed;$$

in the second case, where m_F and m_Q are both negative, we replace m_F by its lower bound and take the absolute value, and replace m_Q with its lower bound from Lemma 3.2.2 and again take the absolute value to obtain

$$\frac{\ell(i) + y_*}{i - x_0} - \left(\frac{\ell(i)}{i - x_0} - E \right) > 2Ed.$$

Either of these two inequalities implies

$$\frac{|y_*|}{i - x_0} + E > 2Ed$$

which in turn implies

$$\frac{4}{3}|y_*| + E > 2Ed,$$

since $i - x_0 > \frac{3}{4}$. The positive integer d on the right side of this inequality is also part of the definition of y_*, and it would be convenient to eliminate it from the inequality. To this end, consider what happens to the closest possible point of intersection as d varies, and take specifically the second case, where m_F and m_Q are both negative. We have two lines, one that contains the points (x_0, y_i) and (i, b) for $b > -\ell(i)$, and another that contains the points $(x_0, y_i + \frac{d}{iq})$ and $(i, -\ell(i))$. As d is increased, the point of intersection moves further and further away from (x_0, y_i). The same is true in the first case, where m_F and m_Q are both positive—more so, in fact; if d is large enough, then the two lines do not intersect anywhere within the polygon $P(N)$. We may therefore without loss of generality assume that d is as small as possible, or in other words $d = 1$. The above inequality then reduces to

$$\frac{4}{3}\left|y_i + \frac{1}{iq}\right| > E$$

which in turn gives us

$$\frac{4}{3}\frac{|c|+1}{q} + \frac{4}{3}\frac{1}{iq} > q + \frac{4}{3}\frac{|c|+1}{q}$$

or

$$4 > 3iq^2.$$

Since $1 \leq w < 2^\beta < \frac{1}{4}q$ by assumption, this cannot possibly be true, and so by contradiction F must not cross through the interior of $Q(N, i)$. It follows that $C(N)$ is indeed exactly the polygon $\bigcap_j Q(N, j)$, which completes the proof that

$$\lim_{N\to\infty} \widehat{C}(N) = D.$$

□

We are now justified in having assumed the truth of these two technical lemmata in proving the main results of the chapter: that chambers of lattice polygons can have any number of edges, and that such chambers can assume (approximately) any convex, centrally-symmetric shape.

Chapter 4

Conclusion

4.1 Questions

We have investigated two sorts of objects of discrete geometry, and found in each case a richer structure than had been supposed to exist. A number of questions remain, however, in some cases suggested by the proofs themselves.

4.1.1 Extendable shellability of uniform matroid complexes

In terms of overall size, the matroid complex we have shown not to be extendably shellable is about as simple as possible while not being trivial. In terms of structure, however, a uniform matroid is simpler: within a ground set $\mathcal{M}$ of size n, declare a subset to be independent if and only if has cardinality k or less. The independence complex of this matroid is also called the k-skeleton of an n-simplex, and it remains an open question whether all such complexes are extendably shellable, as was conjectured [Sim94] by Simon.

4.1.2 First failure of extendable shellability

We have shown that the cross polytope is not extendably shellable in dimension 12, which by suspension also shows it not to be extendably shellable in all higher dimensions. Although 12 is the lowest dimension for which our particular techniques appear to work, it remains an open question which is the lowest dimension for which the cross polytope is not extendably shellable. The present proof suggests a close connection between this question and the question of the smallest oriented uniform matroid with a mutation-free element, although it is possible that a small pseudoplane arrangement with no removable region could exist which did not correspond to an oriented matroid; the definitions are not quite identical.

4.1.3 Mutation-free oriented matroids

We have exhibited a novel method for constructing an oriented uniform matroid with a mutation-free element. It has been conjectured that every orientation of a uniform matroid must contain some mutation; perhaps some variant of the present technique could be applied to that question as well.

4.1.4 Ultimate shellability

Christian Haase defines a complex to be *ultimately shellable* if it and each of its partial shellings is extendably shellable. He has a proof (unpublished) that the suspension of an ultimately shellable complex is extendably shellable, but in work together we were never able to show that the suspension of an ultimately shellable complex is ultimately shellable—with good reason, as it turns out, because the cross polytope is an iterated suspension of an ultimately shellable complex. The present work thus shows implicity that the cross polytope is not ultimately shellable in dimension 11, raising the question of the dimension (probably much lower than 11) for which that property first fails.

4.1.5 Lattice polygons with few sides

The convexity of an approximated pseudochamber implies convexity of the lattice polygon strictly, which possibly explains why computer searches of larger and larger lattice rectangles P have not yet produced examples with complexity greater than 15—the $4n$-gons that arise via the present technique require the function u to be *strictly* concave. Indeed, it is not unreasonable to ask whether there exists a uniform bound c (a foolhardy conjecture would be $c = 4$) such that, if P is any lattice polygon with n edges, the complexity of P is less than or equal to cn.

4.1.6 Linear bounds on complexity

We have found lattice polygons of width $2n-1$ and complexity at least $4n$. This suggests, if these chambers are indeed the largest of the complex, a linear bound on the complexity of a lattice polygon in terms of its smallest dimension, rather than the quadratic bounds which are the best that otherwise exist.

The present construction also depends on making the lattice polygon extremely tall. It may be interesting to investigate what can be done if this dimension is also constrained, or at least to calculate explicitly the size of the lattice polygons generated by the current construction.

4.1.7 Symmetry of chambers

The chambers we have constructed are always approximately symmetric. It is not clear whether this must always be the case for a chamber with many edges, and in particular whether any shape can be approximated regardless of its symmetry.

4.1.8 Position of chambers

The chambers of high complexity we have found are all placed quite close to the edge of the polygon. In the interior of the lattice polygon, there is a wider variety of fissures available, but they seem usually to interfere with each other. If these fissures could all be made to coöperate in some way, it may be possible for such chambers to have more edges than anything occuring near the outskirts of the polygon.

4.2 Acknowledgements

Christian Haase, as mentioned above, first introduced me to the cross polytope and the question of extendable shellability. I am grateful to him for this, and for his generosity in inviting me to take part in his research on the subject.

To Ernest Croot, III, I owe thanks for introducing me to the question of lattice polygon complexity; he was able to find certain bounds (still unpublished) on the size of chambers depending on where they are located within a lattice rectangle. I am also grateful, of course, to Bernd Sturmfels and to Rekha Thomas for originating the problem, and to Diane Maclagan and her coäuthors, firstly for reducing it to a problem I was able to understand, and secondly for popularizing it to the point that it reached my ears. Bernd in particular merits thanks for repeated reminders of the importance of putting the results into writing, without which this dissertation would still not exist in its current form. He also provided much help and advice along the way.

The combinatorics seminar at UC Berkeley deserves some sort of award for luring me away from my chosen field to the point of writing a dissertation on two problems I first learned about there. Thanks go to the various organizers and speakers of that recurring event.

Günter Ziegler read an early draft of this manuscript and made many helpful comments and corrections, several of which have made their way into this text—and all of which perhaps should have.

Laurent Bartholdi may not have directly influenced the research of this dissertation, but he has provided me with one very good impetus to get it written: a postdoctoral position in Lausanne, Switzerland.

Obvious thanks go to my advisor Robion Kirby, whose has continued to be a great support to the research and writing of this dissertation even as it has strayed far from his own field of specialty.

Thanks finally to an indispensible network of family and friends, whose unflagging support and sage encouragement reflect a goodness and faith I aspire to emulate.

Bibliography

[BE94] Anders Björner and Kimmo Eriksson. Extendable shellability for rank 3 matroid complexes. *Discrete Math.*, 132(1-3):373–376, 1994.

[BM71] H. Bruggesser and P. Mani. Shellable decompositions of cells and spheres. *Math. Scand.*, 29:197–205 (1972), 1971.

[BR01] Jürgen Bokowski and Holger Rohlfs. On a mutation problem for oriented matroids. *European J. Combin.*, 22(5):617–626, 2001. Combinatorial geometries (Luminy, 1999).

[DK78a] Gopal Danaraj and Victor Klee. A representation of 2-dimensional pseudomanifolds and its use in the design of a linear-time shelling algorithm. *Ann. Discrete Math.*, 2:53–63, 1978. Algorithmic aspects of combinatorics (Conf., Vancouver Island, B.C., 1976).

[DK78b] Gopal Danaraj and Victor Klee. Which spheres are shellable? *Ann. Discrete Math.*, 2:33–52, 1978. Algorithmic aspects of combinatorics (Conf., Vancouver Island, B.C., 1976).

[HMS04] Serkan Hoşten, Diane Maclagan, and Bernd Sturmfels. Supernormal vector configurations. *J. Algebraic Combin.*, 19(3):297–313, 2004.

[HZ00] Masahiro Hachimori and Günter M. Ziegler. Decompositons of simplicial balls and spheres with knots consisting of few edges. *Math. Z.*, 235(1):159–171, 2000.

[McM70] P. McMullen. The maximum numbers of faces of a convex polytope. *Mathematika*, 17:179–184, 1970.

[MT03] Sonoko Moriyama and Fumihiko Takeuchi. Incremental construction properties in dimension two—shellability, extendable shellability and vertex decomposability. *Discrete Math.*, 263(1-3):295–296, 2003.

[PB80] J. Scott Provan and Louis J. Billera. Decompositions of simplicial complexes related to diameters of convex polyhedra. *Math. Oper. Res.*, 5(4):576–594, 1980.

[RG93] Jürgen Richter-Gebert. Oriented matroids with few mutations. *Discrete Comput. Geom.*, 10(3):251–269, 1993.

[Sch50] Ludwig Schläfli. Theorie der vielfachen kontinuität. In *Gesammelte mathematische Abhandlungen. Band I*, pages 168–392. Verlag Birkhäuser, Basel, 1950.

[Sim94] Robert Samuel Simon. Combinatorial properties of “cleanness”. *J. Algebra*, 167(2):361–388, 1994.

[ST97] Bernd Sturmfels and Rekha R. Thomas. Variation of cost functions in integer programming. *Math. Programming*, 77(3, Ser. A):357–387, 1997.

[Zie98] G. M. Ziegler. Shelling polyhedral 3-balls and 4-polytopes. *Discrete Comput. Geom.*, 19(2):159–174, 1998.