DISS. ETH No. 14672, 2002

Complexity of Triangulation

A dissertation submitted to the
Swiss Federal Institute of Technology, ETH Zürich
for the degree of Doctor of Technical Sciences

presented by
Alexander Below
M.Sc. Operations Research, Cornell University, U.S.A.
born 19 June 1974 in Berlin, Germany

accepted on the recommendation of
Prof. Dr. Dr. Jürgen Richter-Gebert, TU München, examiner
Prof. Dr. Peter Gritzmann, TU München, co–examiner
Prof. Dr. Emo Welzl, ETH Zürich, co–examiner

Abstract

In this thesis we consider the complexity of finding subdivisions of polytopes and polytope complexes with certain extremal properties. Subdivisions will include dissections (subdivisions into simplices with vertices among the polytope vertices such that no two simplices intersect in their relative interiors) and triangulations (dissections that form a simplicial complex). Extremal will always mean having a minimal or maximal number of full-dimensional simplices.

In the first part we consider minimal subdivisions. The main result is that it is NP-hard to find the minimal triangulation of 3-dimensional polytopes. We also investigate the relations of minimal triangulations, minimal dissections and minimal triangulations using additional interior points in this context.

In the second part we consider maximal triangulations. The problem of finding a maximal boundary triangulation over all realizations of a polytope, i.e. of all polytopes having the same combinatorial face structure, will turn out very hard, as hard as solving systems of polynomial equations and inequalities — at least NP-hard.

Using the same techniques we were also able to find interesting results about the realization spaces of polytopes: There are 4-polytopes any realization of which has a certain polygon as a face, and the shape of this polygon is prescribed up to projective equivalence. We will show that this result is best possible in some ways and extend it to higher dimensions.

Zusammenfassung

Diese Arbeit behandelt die algorithmische Komplexität des Problems, gewisse extremale Unterteilungen von Polytopen und Polytopkomplexen zu finden. Die Unterteilungen umfassen simpliziale Zerlegungen (Unterteilungen der Polytope in Simplizes, deren Ecken auch Ecken der Polytope sind, sodass sich nie zwei Simplizes im relativen Inneren schneiden) und Triangulierungen (simpliziale Zerlegungen, die einen Simplizialkomplex bilden). Extremal heisst für uns eine minimale oder maximal Anzahl von volldimensionalen Simplizes.

Der erste Teil ist minimalen Zerlegungen gewidmet. Als Hauptergebnis zeigen wir, dass es NP-schwer ist, die minimal Triangulierung eines 3-dimensionalen Polytops zu finden. Ausserdem untersuchen wir die damit zusammenhängenden Beziehungen von minimalen Triangulierungen, minimalen simplizialen Zerlegungen und minimalen Triangulierungen, die zusätzliche innere Punkte verwenden. Dabei wird klar werden, wie stark die Annahme eines Simplizialkomplexes in diesem Zusammenhang ist.

Im zweiten Teil beschäftigen wir uns mit maximalen Triangulierungen. Das Problem, maximale Triangulierungen des Polytoprands über alle Realisierungen zu finden, d.h. wenn wir alle Polytope mit derselben kombinatorischen Seitenstruktur betrachten, stellt sich als sehr schwer heraus, so schwer, wie es ist ein System von Polynomgleichungen und -ungleichungen zu lösen — mindestens NP-schwer.

Wir benutzen die gleichen Techniken dazu, weitere Fragen in der Theorie der Realisationsräume von Polytopen zu beantworten. Zum Beispiel gibt es 4-dimensionale Polytope, die in jeder Realisierung ein gewisses Polygon als Seitenfläche haben, und diese Seitenfläche wird in jeder Realisierung bis auf projektive Äquivalenz die gleiche Form haben. Wir werden zeigen, dass dieses Resultat auf gewisse Art bestmöglich ist, und es auf höhere Dimensionen verallgemeinern.

Acknowledgements

First of all I want to thank my parents for always being there for me, supporting me, (almost) all the time believing in me, and forgiving me for being busy with my thesis and not being as attentive to their issues as I should have.

I do not know where to begin to thank my advisor Prof. Jürgen Richter-Gebert. Jürgen is an extraordinary person. By always showing me that there is an alternative way of looking at things he humbled me and inspired an awe before his superior flexibility of mind. He taught me what many like to but still do not believe, that true enthusiasm does not stop before the limitations life seems to be so full of. All this he achieved not by giving me constant advise, but by living the way he does and only once in a long while keeping me on my track when I was in danger of dwelling on side-issues. All this paired with no hesitation to help whenever he can, in person, over the phone, by email. He claims he has a bad conscience for not having time for anybody; this is nonsense as far as I am concerned and always has been.

Emo Welzl and Peter Gritzmann as my co-examiners have been very helpful with their questions and comments regarding the thesis.

Jesús De Loera was the person who cunningly got me into this whole "game of triangulations". By asking me innocent questions that we trivial sometimes and very hard at other times he almost tricked me into getting involved in the topic which later was to become my thesis. I will not forget our eternal (sic!) discussions about the wording of the first paper he and I were writing. I learnt so much and am equally eternally grateful.

I thank Prof. Günter Ziegler for being an attentive listener and insight-giver during the past years whenever I was in town and needed help. He was the one who steered me into the field of discrete mathematics. Martin Grötschel ignited my first interest in discrete math in his lectures on linear and combinatorial optimization and during a seminar was responsible for my meeting with Jürgen Richter-Gebert. He was also the one that opened for me the path to Cornell, my last station before coming to the ETH.

My flatmate Igor Režan was the try-out audience every time I almost had a nervous breakdown before the next talk. I had so much fun talking with you, being mad at you, being embarrassed by your singing in public ("Noooon, rien de rieeeeen..."), and – not to forget – watching Becker with you.

My friends Uli Wagner, Falk Tschirschnitz, Ingo Schurr, Ulli Kortenkamp, Matthias John, Jochen Giesen, Udo Adamy, Mark Cieliebak, Zsuzsa Liptak, Julian Pfeiffle, Hans-Martin Will, and Börnie Seybold also gave me so many helpful hints when I asked for an ear or a mind, mathematically or about life in general. They were also rebuilding my self-confidence when just nothing seemed to go the right way. Ulli basically showed me how to use a computer.

I thank my friends Enno Brehm and Olli Bröker for putting things into a lighter, funnier perspective, sometimes in the most absurd way, when I was so tense about this "math thing" and could not seem to relax.

My friends Michael Schmidt and Kay "Mutti" Neuenhofen whom I have not seen for some years now, I thank you for being so cool and being there in my mind when sometimes I was confused and was asking myself "what would Kay/Michael do now?"

My past office mates Artur Andrzejak, Beat Trachsler, Ulli Kortenkamp, Vanessa Krummeck, Czaba Tóth, and my current office mate Miloš Stojaković were most patient with my ramblings, outbursts or snubbing quiet moments.

Eventually, I want to thank my girlfriend Fabia Denninger for being there (during the week over the phone, over the weekend in person) when I needed to talk and for being gone when I needed to sit down and work and not see anybody. Fabia, you are the best girlfriend imaginable.

A.B., Zürich, September 2002

Contents

. . . so I figured out what you gotta do, man, every time you're lookin' for a piece of *action* and you ain't gettin' that, man, you know what you gotta do, baby, you better *try harder* . . .

Janis Joplin

Chapter 1

Introduction

A general paradigm in algorithmic mathematics is the decomposition of complex structures into smaller pieces – often it is easier to handle many small objects than one big one. Think of numerical integration where shapes are approximated by a collection of cubes, of topology where topological spaces are broken down into cell complexes in order to compute topological invariants, or of algebra where exact sequences are applied to get a computational handle on groups or other structures.

The subdivisions of polytopes into simplices is an incarnation of this paradigm. We study the following type of subdivisions:

Definition 1.1 *A* triangulation *of a d-dimensional polytope P is a collection T of d-dimensional simplices such that*

1. *the vertices of the simplices in T are among the vertices of P,*
2. *the union of the simplices in T equals P, and*
3. *the intersection of any two simplices in T is a face of both simplices (which is possibly empty).*

The size *of a triangulation is the number of d-simplices it contains.*

Observe that, if the dimension is at least 3, polytopes can have triangulations with different sizes (see Figure 1.1).

In this thesis we study *minimal* and *maximal* triangulations, i.e. triangulations whose size is minimal or maximal among all triangulations of the polytope.

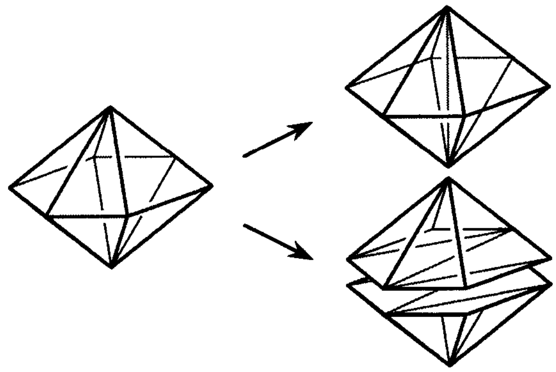

Figure 1.1: *Two triangulations of the bipyramid over the hexagon having sizes* 6 *and* 8*, respectively*

1.1 Applications of Optimal Triangulations

We give now a brief list of applications of minimal and maximal triangulations in various fields:

- One standard method for the computation of the volume of a polytope uses a triangulation since the volume of a simplex is easily expressed as a determinant. It is conceivable that triangulations with a small number of simplices help in keeping down the computation time of the volume. Note that the volume computation is $\#P$-hard if the dimension is part of the input [24].

- Maybe the first impulse to study minimal triangulations came from the approximation of fixpoints of continuous maps. In economics and game theory market equilibria are computed using fixpoints. Efficient algorithms for approximate fixpoints exploit the combinatorial structure of a suitable triangulation of the domain of the map [59]. These algorithms are very sensitive to the triangulation used since they are incremental and move from simplex to simplex. One measure for the suitability of a triangulation is the number of simplices used. The cube, while not the only interesting example, is particularly interesting: Stacking triangulated cubes may give rise to triangulations of unbounded or large regions in euclidean space. The problem of finding the minimal triangulation of the d-cube has been extensively studied, the answer is only known up to $d = 7$ [34], other than that only non-matching upper and lower bounds are known [19, 33, 41, 56].

- Closely related to this is the area of mesh generation with its applications in finite element methods [9]. Finite element methods have proved indispensable for physical simulation. These methods discretize the simulated domain—for example, the air around a wing—by dividing it into many small "elements," like triangles or tetrahedra. The goal of these methods is then the approximation of solutions for differential equations by only considering local interactions within one simplex or between adjacent simplices in the mesh (the triangulation).

  Also in simplicial meshes suited for finite element methods the triangulation plays an important role in the performance of the method. A small number of simplices will certainly help in keeping the number of computations down. However, besides size of the triangulation other measures are important, for example the concrete shapes of the tetrahedra. There are a number of other measures of a triangulation whose optimization will warrant good performance of the application, these include a small number of "slivers" (long and thin simplices that introduce numerical instabilities), minimizing the largest dihedral (most obtuse) angle (also causing instabilities), minimizing the total length of interior edges etc.

- Understanding minimal triangulations of convex polytopes is also related to the problem of characterizing the f-vectors of triangulations of balls and polytopes (see open problems in [11]). In fact, the study of minimal triangulations of topological balls and of polytopes found an application in the calculation of maximum rotation distance of binary trees. Sleator et al. [55] use a nice volume argument in hyperbolic space which links the size of a minimal triangulation via a universal lower bound on the hyperbolic volume of a simplex to the overall volume of special polyhedra which turns out to be linear in the number of vertices.

- Triangulations also play an important role in the classification of real algebraic curves. In fact, by a powerful combinatorial technique due to Viro ([35], see also [50]) it is possible to construct real algebraic curves, so-called T-curves, whose topological shape is determined by a triangulation of the integer points in a given polytope. Maximal triangulations correspond in this framework to particularly interesting topological types.

1.2 The Results in this Thesis

One of the main goals in our research was to determine the complexity of finding minimal or maximal triangulations of convex polytopes. The computational geometry literature has several articles dealing with polynomial algorithms for finding triangulations of optimal size [2, 28]. In particular, Bern and Eppstein asked in 1992 whether there is a polynomial-time algorithm to compute a minimal triangulation of a 3-polytope (Open Problem 12 in Section 3.2 [9]). We will prove that such an algorithm cannot exist unless $P = NP$.

Extremal Triangulations are not Invariants of the Face Lattice

The first question which was anwered in this context was whether the minimal or maximal triangulations are invariants of the face lattice of a 3-polytope, i.e. if different coordinatizations might have different minimal or maximal triangulations. In dimension 2 all convex n-gons have only triangulations of size $n-2$ and all of them are present in all coordinatizations. Two polytopes with the same oriented matroid (i.e. with the property that all simplices spanned by $d+1$ vertices have the same orientation) have the same triangulations [12]. This question is therefore a reformulation of the question whether the extremal triangulation solely depends on the face lattice or whether *interior* oriented matroid information is necessary. This question is of course interesting when trying to design algorithms to *find* small triangulations.

In three dimensions this was known for maximal triangulation: the regular 3-dimensional cube has a maximal triangulation of size 6 while a perturbation gets it up to 7 tetrahedra (see Chapter 4).

For minimal triangulations Richter-Gebert and independently Brehm solved this problem by supplying the polytope in Figure 1.2: This polytope has two vertex-edge chains on the boundary, i.e. sequences of vertices in which consecutive vertices are joined by edges, such that all vertices of one chain are joined to the two end vertices of the other and vice versa. These vertex-edge chains all of whose vertices are joined to exactly two other vertices have a very powerful property if there are only few vertices outside the vertex-edge chain: *Any small triangulation must use the diagonal connecting the two vertices.* If the two pairs of end points are coplanar this is not possible, of course. In Chapter 2 we will show that in this case the minimal triangulation has 10 tetrahedra, but if the vertex-edge chains are moved a bit apart (such that their convex hulls do not meet) then both diagonals can be used and a minimal triangulation uses 9 tetrahedra.

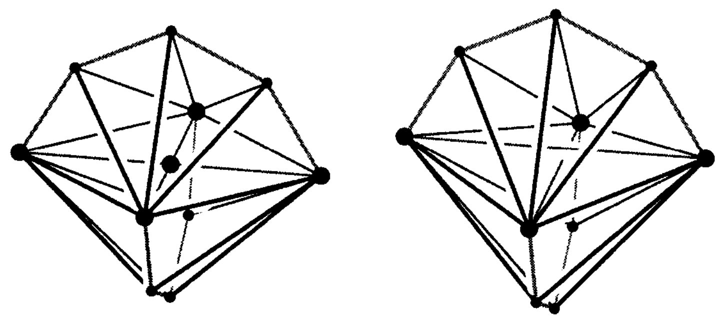

Figure 1.2: *Two coordinatizations of the 7-cabbage which have minimal triangulations of sizes 10 and 9*

Theorem 1.2 *The minimal size of a triangulation of a convex 3–polytope is* not *an invariant of the face lattice:*

1. *There is a simplicial convex 3-polytope with 10 vertices for which the minimal number of tetrahedra possible in a triangulation depends on its coordinates.*
2. *The example is smallest possible in dimension and number of vertices.*

The proof appeared in [5], we present it in Chapter 2.

Minimal Dissections and Triangulations using Interior Points

The power of the vertex-edge chains spawned the solution of other problems concerning minimal subdivisions generalizing triangulations. Relaxing the face-to-face condition in the definition of triangulations we get the notion of *dissections*: a subdivision of a polytope such that any two simplices must not meet in their interiors, but not necessarily in common faces (see Figure 1.3).

A natural question was whether dissections could have smaller size than minimal triangulations. Independently Böhm [14] and Gritzmann and Klee [32] raised this issue. As pointed out in Section 8.4 of [32], this question is relevant in the study of complexity classes of basic problems in computational convexity. The answer is yes:

Theorem 1.3 *There is a family of polytopes P_m^{diss} (m an integer parameter) allowing a dissection which is smaller than every triangulation of P_m^{diss}. The gap is linear in the number of vertices.*

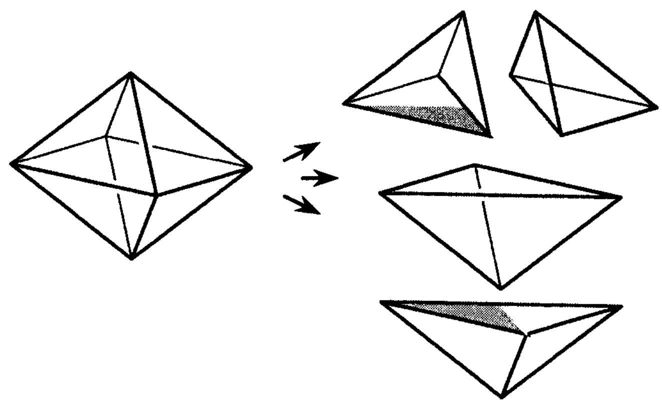

Figure 1.3: *A dissection of the octahedron which is not a triangulation*

Also we can relax the notion of triangulation by allowing the simplices having other vertices than just the polytope vertices, namely additional interior points. These are often called Steiner points. We refrain from this notion as it is confusing with the original meaning of Steiner point (which refers to the additional points in a special problem, namely finding a spanning tree of minimal total length of points in the plane using additional vertices).

Could a triangulation using additional interior points have fewer tetrahedra than a triangulation without interior points? This question was posed several times [14, 20, 32, 34]. It reminds of the fact that Delaunay triangulations of point sets can have smaller complexity when throwing in more points. (Notice that Delaunay triangulations are almost never minimal triangulations. We come back to this issue when we talk about approximations.) Again, the answer is yes.

Theorem 1.4 *1. There is a family of polytopes P_m^{int} allowing a triangulation using an interior point which is smaller than every triangulation of P_m^{int} not using interior points. This gap is linear in the number of vertices.*

2. Even more, the gap and the minimal number of used interior vertices can be prescribed: Given two numbers $h \geq 1$, $k \geq 1$ there is a simplicial convex 3–polytope P such that every triangulation of P using less than h interior points has at least k tetrahedra more than a triangulation of P with h suitably chosen interior points.

The proof which appeared in [5] can be found in Chapter 2.

Finding Minimal Triangulations is NP-hard

Let us make a slight digression to non-convex polytopes: For non-convex polyhedra a triangulation cannot always be found. As an example we note the *Schönhardt polytope* [9, 42, 51, 52] which will play an important role in the sequel. It is obtained by twisting the top face of a triangular prism in a clockwise direction (see the first transformation in Figure 1.4). Ruppert and Seidel [51] showed that it is NP-hard to decide if a polyhedron can be triangulated. (Note however that every polyhedron can be triangulated if one allows the simplices to have vertices other than the polytope vertices [9, 16, 17].)

Now consider vertex-edge chains patched to the sides of the Schönhardt polytope; we call the resulting convex polytope *cupola* (see the second transformation in Figure 1.4). We can construct polytopes from this polytope which

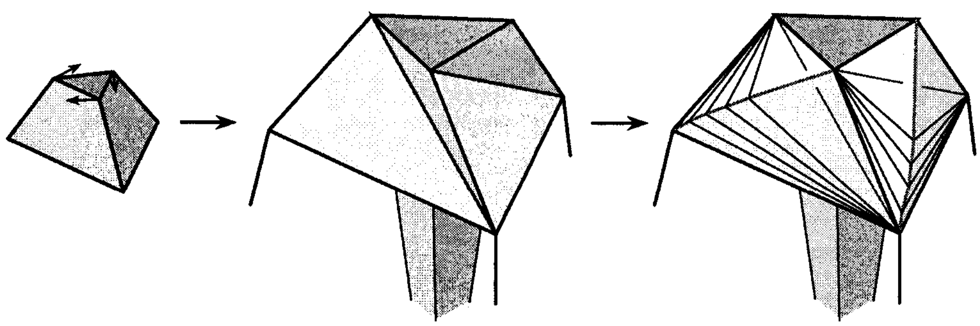

Figure 1.4: *Patching the sides of a Schönhardt polytope with vertex-edge chains we obtain the cupola, glued to a bigger polytope it has the property that the vertex triangulating the top face must lie in a special cone—the visibility cone*

have a remarkable property in every minimal triangulation: The fourth point of the tetrahedron which triangulates the top face has to lie in a special cone, the *visibility cone*. Furthermore, we are able to prescribe the visibility cones at will. Also we can glue many cupolas to one frame polytope and get many restrictions of points triangulating some boundary triangles lying in specified visibility cones.

These observations turn out to be the decisive ingredients in the proof of NP-hardness of finding the minimal triangulation of 3-polytopes. Define the following decision problem:

MINTRIANG(d)	
Given:	A d-dimesional polytope P and a natural number K
Question:	Is there triangulation of P with less than K simplices?

Theorem 1.5 MINTRIANG(3) *is an NP-complete problem.*

We prove NP-completeness by giving a polynomial transformation from the well-known Satisfiability problem:

SATISFIABILITY (SAT)	
Given:	A logical formula f consisting of c logical clauses over v variables
Question:	Is there an assignment of true/false to the variables which satisfies all clauses?

Given such a logical formula f we will construct a 3-dimensional polytope, our so-called *logical polytope* P_f and a number K_f which will admit a triangulation using less than K_f tetrahedra if and only if the formula f admits a satisfying truth assignment. This construction will be such that the encoding length of the vertex coordinates of the polytope are bounded by a polynomial in c and v. The proof appeared in [6]. We give it in Chapter 3.

The corresponding decision problems for minimal dissections and minimal triangulations are:

MINDISSECT(d)	
Given:	A d-dimesional polytope P and a natural number K
Question:	Is there a dissection of P with less than K simplices?

MINTRIANGIP(d, l)	
Given:	A d-dimesional polytope P and a natural number K
Question:	Is there a triangulation of P using at most l interior points with less than K simplices?

Since the vertices of logical polytopes will be in general position, it does not have dissections which are not triangulations, hence MINDISSECT(3) is also NP-complete. By a supplementary construction on top of the logical polytope we will also be able to show the result for triangulations using a specified maximum number l of interior points, MINTRIANGIP($3, l$) is NP-complete if l is not part of the input.

There are other interesting conclusions from this theorem. First of all, the hardness of the decision problem implies the hardness of the optimization problem: If there was a polynomial algorithm which finds the minimal triangulation (dissection) of the polytope P_f then we could also determine whether this triangulation (dissection) is smaller than the number K_f. This proves that it is NP-hard to find the minimal triangulation (dissection) of a polytope.

Second, these completeness and hardness results can be extended to polytopes of any fixed dimension ≥ 3: Taking the pyramid over a polytope P results in a polytope P_{pyr} of one dimension higher. This polytope has the property that each triangulation (dissection) consists of simplices which are pyramids over simplices in the original polytope P. Moreover, these lower-dimensional simplices form a triangulation (dissection) of P. Hence the triangulations (dissections) of P are in a size-preserving one-to-one correspondence with the triangulations (dissections) of P. Thus, it is at least as hard to find the size of the smallest triangulations in four dimensions as it is in three. Repeating this pyramid construction a sufficient number of times, adding a new dimension each time, we have the following corollaries (the third result was recently obtained in [45] via a direct transformation of 3-SAT):

Corollary 1.6 *1.* MINTRIANG(d), MINDISSECT(d), *and* MINTRIANG(d, l) *are NP-complete for all $d \geq 3$ and $l \geq 0$ (d and l are not part of the input).*

2. Finding the minimal triangulations (dissections) of polytopes in fixed dimension d is NP-hard for all $d \geq 3$.

3. Finding the minimal triangulations (dissections) of the boundaries of polytopes in fixed dimension d is NP-hard for all $d \geq 4$.

4. Fix a number l. Finding the minimal triangulation of 3-dimensional polytopes using at most l interior points is NP-hard.

The proof can be found in Chapter 3.

Notice that Corollary 1.6.4 is not really what one wants: It would be interesting to know whether finding the minimal triangulation using an arbitrary number of interior points is also hard. We strongly conjecture this. This is an open problem.

The result also generalizes from realized polytopes to *matroid polytopes* [12]. This is so since every polytope gives rise to a matroid polytope (computable in polynomial time) and the triangulations of the two are in 1-to-1 correspondence.

Corollary 1.7 *For dimensions 3 and higher computing a minimal triangulation of a matroid polytope is NP-hard*

Once the hardness of a problem is established the natural question is for approximations. First notice that Delaunay triangulations are a good approximation since they have quadratic worst-case complexity [30] which as bad as it gets: the upper bound of the size of a triangulation is $O(n^{\lfloor (d+1)/2 \rfloor})$ [49]. But the so-called *pulling triangulation* [39] is an easy 2-approximation can be computed in time $n \log n$. It is obtained by singling out one vertex p of the polytope, triangulating all facets not containing p. Then the tetrahedra obtained as the convex hull of the triangles and p provide a triangulation. The size of the triangulation is bounded by the size of the boundary triangulation which is a complete graph and therefore $2n - 4$. On the other hand each triangulation has at least size $n - 3$ [49]. Recently, Chin et al. [18] were able to improve the approximation ratio to $O(2 - 1/\sqrt{n})$.

There are however classes of polytopes for which it is easy to give minimal triangulations. One is the class of *stacked polytopes*. These are polytopes that can be obtained from a d-simplex by attaching one d-simplex after the other to it, always maintaining convexity [5, 49]. Recently [60] this class was somewhat enlarged using vertex-edge chains.

Note that the question of finding maximal triangulations is only partly answered. If the dimension part of the input the problem was recently shown to be $\#P$-hard [24].

Extremal Triangulations over all Realizations

Since minimal and maximal triangulations are not invariants of the face lattice it is natural to ask for the minimal or maximal triangulations among all realizations of a polytope. (A realization is a polytope with the same face lattice.) In two dimensions this is again easy since all triangulations are present in all realizations of a polygon. In three dimensions this is basically open.

What makes this problem intrinsically different is that the face lattice does not determine the oriented matroid of the polytope vertices. The oriented matroid encapsulates exactly the combinatorial information which is necessary to decide if a collection of simplices is a triangulations and which is not. This seems to suggest that in order to compute extremal triangulations an algorithm has to look at all realizations of the polytope.

In three dimensions there might be some hope: The spaces of all realizations of a polytope is simply connected, this is a consequence of the proof Steinitz' theorem which gives a combinatorial characterization of all face lattices of convex 3-polytopes [61]. The proof is constructive, so it is possible to find at least one realization of a given face lattice. Another approach to find a realization uses "stress graphs" [46].

On the other hand, Richter-Gebert [46] has shown that the space of all realizations of a 4-polytope can be arbitrary complicated, as complicated as a solution set of any polynomial equation and inequality system. For this *universality theorem for polytopes* he proves that the problem of finding a realization of a face lattice is at least as hard as the problem of finding a solution to a system of polynomial equations and inequalities, the *existential theory of the reals* (ETR).

EXISTENTIAL THEORY OF THE REALS (ETR)	
Given:	A system of polynomials $\{f_j\}$ and $\{g_k\}$ with algebraic coefficients in variables $x_1, \ldots, x_n$
Question:	Is there an assignment of real numbers to the variables x_i which satisfies all equations $f_j(x_1, \ldots, x_n) = 0$ and all inequalities $g_k(x_1, \ldots, x_n) < 0$?

It is unknown whether the existential theory of the reals is in NP. It is in $PSPACE$ [15], but $PSPACE$-hardness is also not known. However, it is easy to see that ETR is at least NP-hard (simple transformation from the SAT problem).

We will use the ETR to show that the following problem is hard:

MAXTRIANG∂FL(d)	
Given:	The face lattice of a d-dimensional polytope P and a number K
Question:	Is there a realization of P which admits a boundary triangulation using more than K d-simplices?

Theorem 1.8 MAXTRIANG∂FL(5) *is as hard as the existential theory of the reals, hence at least NP-hard.*

We use Richter-Gebert's universality construction to encode a 5-polytope most of whose facets are pyramids over pyramids over polygons. These 4-polytopes have the same triangulations in every realization, so it is easy to triangulate them maximally. (The triangulations of a pyramid are in 1-to-1 correspondence with the triangulations of the ground face. Hence all triangulations of a pyramid over a pyramid over a polygon are present in all realizations.)

However, in our construction one facet is a pyramid over a hexagonal prism. Our universality construction ensures that only two realizations of this prism are possible, one with a large maximal triangulation (17 tetrahedra) and another with a smaller maximal triangulation (14 tetrahedra). Also a given instance of the ETR (a so-called *driving* system of polynomial equations and inequalities) is encoded in the polytope such that the large triangulation can be attained if and only if the instance has a solution. So by finding the maximal boundary triangulation over all realizations it is possible to decide whether the driving ETR instance has a solution.

This result is interesting in two ways: First, it is the first triangulation result which takes into account the space of all realizations of a polytope. Second, it is one of the first applications of the universaliy theorem for polytopes.

Naturally, finding minimal triangulations over all realization is also interesting. We strongly conjecture that they are very hard, as well. Unfortunately, due to the reliance of our proof on maximal triangulations of prisms this construction does not just carry over to minimal triangulation: all realizations of a prism over a polygon have the same minimal triangulations [25].

Prescribing Exact Shapes of Faces of Polytopes by their Face Lattice

In the years before the universality theorem for polytopes several researchers had found examples of polytopes whose face lattices put restrictions on the shape of some of their faces. Not all realizations of these faces could be completed to a realization of the big polytope. The universality theorem also generates instances of this category of polytopes that prescribe some property to a face: in every realization the slopes of the edges of a polygonal face encode a solution to a system of polynomial equations and inequalities.

The question now was whether it is possible to prescribe the exact shape of a face in all realizations. With exact shape we mean that in every realization the face is determined up to a projective transformation. This is all one can ask for since projective transformations of a polytopes induce projective transformation of all of their faces and leave the face lattice invariant. We can give an affirmative answer for faces in dimensions $d \geq 2$:

Theorem 1.9 *1. Let G be a d-dimensional (realized) polytope with algebraic vertex coordinates. Then we can construct a $d+2$-dimensional polytope P which has a face G which in every realization is the image of a projective transformation of G.*

This result is especially interesting since it presents a nice interface for further results in this area of realization spaces in view of that the universality theorem itself always seems kind of hard to handle as a basic building block.

We conjecture that d-faces of $d + 1$-polytopes can not be prescribed. While in dimension $d = 2$ this is a consequence of Steinitz' theorem, in higher dimensions this is open.

1.3 Organization of the Thesis

In the first part of this thesis we talk about minimal triangulations (and dissections) of polytopes. In Chapter 2 we give the formal definitions of triangulations and dissections, supply the vertex-edge chain lemma and give application of this lemma. The results are mainly due to Richter-Gebert and De Loera and independently Brehm. However, they present nice applications of the key lemma (the proof of which is supplied by the author). The results already appeared in a joint article [5]. Chapter 3 is dedicated to the proof that $\textsc{MinTriang}(3)$ and $\textsc{MinTriangIP}(3, l)$ are NP-hard. This construction was mainly done by the author and is to appear in [6].

In the second part we talk about maximal triangulations and their behavior under coordinate change as well as the prescribing of the shapes of faces. The results are joint work of Richter-Gebert and the author. In Chapter 4 we give an introduction to the concept of combinatorial polytopes and their realizations, prescribing of properties, and to Richter-Gebert's universality theorem for polytopes. A short appendix gives the prerequisites of the theory of projective spaces we will use in the later chapters. In Chapter 5 we prove Theorem 1.9. Many of the techniques presented there we will reuse in Chapter 6 when we show that $\textsc{MaxTriang}\partial\mathrm{FL}(5)$ is hard.

Part I

Finding Minimal Triangulations of Polytopes

Chapter 2

Minimal Triangulations and Dissections

2.1 Decomposing Polytopes Into Simplices

In this thesis we will consider three concepts of decomposing d-dimensional convex polytopes into simplices (for an overview see also [39]). A *dissection* of a d-dimsional polytope P is a collection D of d-dimensional simplices such that

1. the vertices of the simplices in D are among the vertices of the polytope P,
2. the union of the simplices in D is the polytope P, and
3. no two simplices in D intersect in their interior.

A *triangulation* T of P is a dissection of P with the additional condition that any two simplices in T intersect in a common (possibly empty) face. One could say equivalently that a triangulation is a dissection such that the simplices and their (iterated) intersections form a simplicial complex. Figure 2.1 shows a dissection of the octahedron. Note that the top left simplex and the bottom front simplex meet in a two-dimensional set which is not a face to either of them. Hence, this dissection is not a triangulation.

A *triangulation* of a d-dimensional polytope P *using interior points* is a collection T of d-dimensional simplices such that

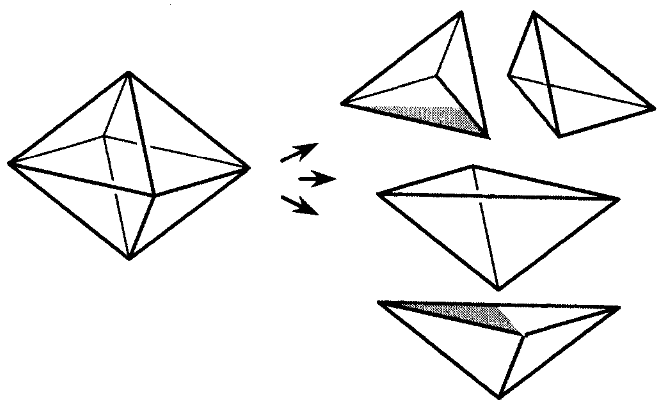

Figure 2.1: *A dissection of the octahedron which is not a triangulation*

1. the vertices of the simplices in T are *contained* in P,

2. any two simplices in T intersect in a common face, and

3. the union of the simplices in T is the polytope P.

The vertices of the simplices in such a triangulation fall of course in two classes: the ones that are also vertices of the polytope and the ones that are in the (relative) interior of (some face of) the polytope. Sometimes we will prescribe or bound the number of these vertices: we will speak of *triangulations using (at most) l interior points*.

In this thesis we are only interested in combinatorial data of these decompositions, more precisely the size: The *size* of a triangulation (dissection / triangulation using interior points) is the number of its (d-dimensional) simplices.

In dimensions greater than two different triangulations and dissections can have different sizes: For instance, the bipyramid over the hexagon admits triangulations with six tetrahedra (join the six edges of the hexagon with the vertical interior edge) and with eight tetrahedra (triangulate the hexagon and join each triangle with the top and with the bottom apex), see Figure 1.1.

We will say that a dissection / triangulation / triangulation using interior points is *minimal / maximal* if its size is minimal / maximal among all dissections / triangulations / triangulations using interior points.

Some of the results in this thesis are (in hindsight) very easy to come by once the main ingredients have been figured out. The proofs to these results present a premium way into the subject: As they build on each other, we can get a

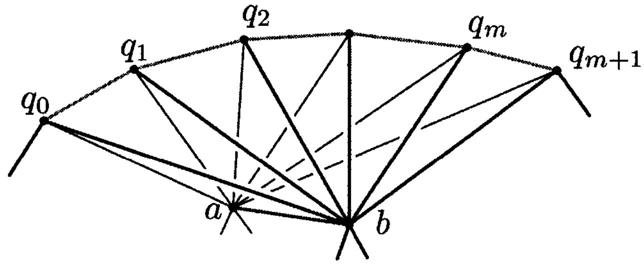

Figure 2.2: *A vertex-edge chain of length* 4

nice view on the key ideas that are also used in the more complex treatments of Chapter 3.

2.2 Vertex-Edge Chain Lemma

We now present the lemma that showcases the main non-trivial effect we will use in all the examples. The lemma shows that a certain substructure in the *face lattice* of a polytope forces certain interior edges to appear in triangulations of small size.

Definition 2.1 *Let P be a not necessarily convex 3-dimensional polytope having the following properties:*

1. *It contains the following collection of triangular facets:* (a, q_i, q_{i+1}), *and* (b, q_i, q_{i+1}) *for* $i = 0, \dots, m$ *(see Fig. 2.2).*

2. *The points* q_i *are in convex position.*

3. *All edges* (q_i, q_j) *for* $|i - j| \geq 2$ *go through the interior of P.*

We say that a polytope satisfying the assumptions 1.–3. has a vertex-edge chain $q_0, \dots, q_{m+1}$ *of* length m *with* diagonal (a, b).

Such a vertex-edge chain has a natural small triangulation consisting of the tetrahedra spanned by the diagonal and consecutive vertices on the chain. We cannot prove that this is a subtriangulation of every minimal triangulation (since this is false). But we can prove that if the diagonal is not present in the triangulation then we have a guaranteed lower bound on the size of the triangulation.

Lemma 2.2 ***Vertex-Edge Chain Lemma*** *Let P be a polytope on n vertices with a vertex edge chain of length m and T a triangulation of P (possibly with interior points) that does not use the diagonal (a, b). Then T uses at least $n + m - 3$ tetrahedra.*

This lemma creates the possibility to force the occurrence of many tetrahedra when some interior edges are absent. The lemma is so useful because its local assumptions (the vertex-edge chain on the boundary and the non-existence of just one interior edge) have global effects (a relatively large triangulation).

Before we come to the proof we state a lemma which links the number of the tetrahedra of a triangulation to the number of interior points and interior edges. It follows that minimal triangulations are the ones using the smallest number of interior edges. In the proof of Lemma 2.2 it will turn out to be much easier counting interior edges than counting tetrahedra.

Lemma 2.3 *Let P be a 3-polytope with n vertices. For a triangulation T of P with n_{int} interior points, the number of tetrahedra in T is related to the number of interior edges e_{int} of T by the formula:*

$$|T| = n - n_{int} + e_{int} - 3. \tag{2.1}$$

Even though an easy consequence of the Euler formula for simplicial complexes, we list the proof for the sake of completeness.

PROOF: The triangulation T defines a simplicial complex. Let e denote the number of edges of T on the boundary of P, t the number of triangles of T in the boundary of P, and t_{int} the number of interior triangles of T. Since P is a topological ball the Euler formula implies:

$$n + n_{int} - e - e_{int} + t + t_{int} - |T| - 1 = 0.$$

Each tetrahedron has four triangular faces, so if we sum over all tetrahedra we count all interior triangles twice, so

$$4|T| - t - 2t_{int} = 0.$$

Using these two equations we can eliminate t_{int}, we obtain

$$n + n_{int} - e - e_{int} + \frac{1}{2}t + |T| - 1 = 0 \tag{2.2}$$

The vertices edges and triangles of T on the boundary of P define a complete planar graph, therefore

$$\begin{aligned} e &= 3n - 6 \\ t &= 2n - 4 \end{aligned}$$

Substituting these two equations into Equation 2.2 gives the desired result. □

Notice that Lemma 2.3 is false for dissections since dissections are not in general simplicial complexes. So in particular, this lemma represents a way to distinguish dissections and triangulations.

PROOF OF LEMMA 2.2: By Lemma 2.3 it suffices to show that T has at least m interior edges. The proof is then complete since n_{int} is a non-negative number.

The proof proceeds by induction on the length m of the vertex-edge chain. The claim is clearly true for $m = 0$. Fix a triangulation T of P.

Case 1: There is a q_i for $1 \leq i \leq m$ which is not incident to an interior edge in T. We now show how to invoke induction: A vertex q_i untouched by an interior edge belongs to the tetrahedra $\sigma_{i,a} = (a, q_{i-1}, q_i, q_{i+1})$ and $\sigma_{i,b} = (b, q_{i-1}, q_i, q_{i+1})$. This is so because the triangle (a, q_i, q_{i+1}) is in some simplex, and if the fourth point is some other vertex besides q_{i-1} or b we have an interior edge touching q_i. Furthermore the fourth point cannot be b since in this case the edge ab would be present. By chopping off these two tetrahedra together with the vertex q_i we can apply induction to guarantee that the remaining triangulation $T \setminus \sigma_{i,a}, \sigma_{i,b}$ has at least $m - 1$ interior edges. Together with the edge $q_{i-1}q_{i+1}$ they account for m interior edges in T. Note

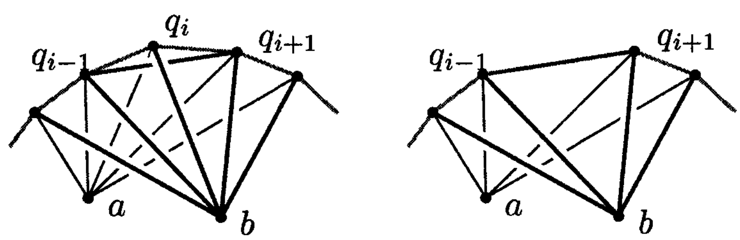

Figure 2.3: *Chopping off vertex q_i which is not incident to any interior edge*

that possibly (q_{i-1}, q_{i+1}) is now a non-convex edge. Invoking induction might not be possible then. We have seen though that we do not use convexity explicitly, but only that all edges incident to the q_i which are not expressly on the boundary are interior edges. This will carry on in Case 2. Hence there is no problem with the induction.

Case 2: All q_i, $1 \leq i \leq m$, are incident to an interior edge in T. Now we do not invoke induction, rather we will show the claim directly. We set up a one-to-one (but not necessarily onto) map from the set $\{q_1, \ldots, q_m\}$ to a subset of the interior edges that touch them: The vertices q_i come along a polygonal curve in a canonical order which is reflected by their indices. We mark or orient the interior edges $q_i v$ that touch a vertex q_i as follows: If $v \notin \{q_0, \ldots, q_{m+1}\}$, we call the edge $q_i v$ *special*, otherwise we orient it from smaller to larger index. For the vertices q_i with special edges incident to them, we map q_i to one of those (see Figure 2.4.a). If a vertex q_i has no special edges, but has outgoing interior edges, we map it to the outgoing edge $q_i q_k$ with the smallest index k (see Figure 2.4.b). We are left with the case of those vertices q_i that have only incoming interior edges incident to q_i. Consider the triangle (a, q_i, q_{i+1}). It has to be in some tetrahedron of T whose fourth point is bound to be a q_{j_a} with $j_a < i$, otherwise q_i enjoys the presence of an incident special or outgoing edge. Likewise (b, q_i, q_{i+1}) is in a tetrahedron with fourth point q_{j_b} with $j_b < i$. If both $j_a = j_b = i-1$, there can be no interior edges incident to q_i (see above), a contradiction to the case assumption. Let j be any of j_a, j_b such that $j < i-1$. Map q_i to $q_j q_{i+1}$. See Figure 2.4.c.

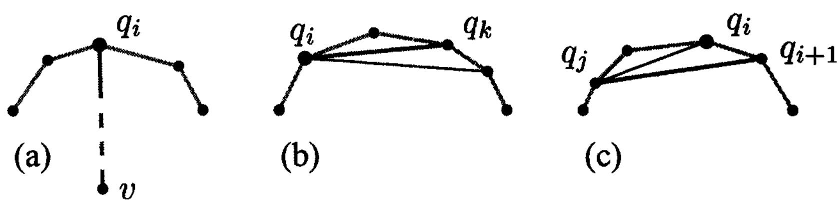

Figure 2.4: *Three ways of mapping vertices to interior edges*

We claim that the given map is one-to-one. If some vertex q_i maps to the special edge $q_j v$, then necessarily $i = j$. There are potentially two vertices that can be mapped to an interior edge $q_j q_k$ with $j < k$: q_j when $q_j q_k$ is the chosen outgoing edge of q_j and q_{k-1}, in case q_{k-1} has only incoming edges. In the latter case one of the tetrahedra (a, q_j, q_{k-1}, q_k) and (b, q_j, q_{k-1}, q_k) has to be in the triangulation, and q_j will be mapped to the smaller indexed edge $q_j q_{k-1}$. This is an interior edge since $j < k - 2$, so q_j cannot also be mapped to $q_j q_k$. The injectivity of the map is proven. $\square$

In order to show the power of the key lemma we use it to give answers to some of the initial questions.

2.3 Consequences of the Vertex-Edge Chain Lemma

2.3.1 Dissections can be Smaller than Minimal Triangulation

Theorem 2.4 *There is an infinite family of polytopes P_m^{diss} on $n = 4 + 2m$ vertices (m an integer parameter) allowing a minimal dissection of size $2m - 2$, but the minimal triangulation of P_m^{diss} has $3m + 1$ tetrahedra.*

Notice that the gap between minimal dissection and minimal triangulation is minimal in the number of vertices.

The construction will rely on the degeneracy of the constructed point configuration – dissections of polytopes with vertices in general position are automatically triangulations.

PROOF: We first show the construction of P_2^{diss} which has a minimal dissection using 6 tetrahedra and a minimal triangulation using 7 tetrahedra. It is obtained as the convex hull of a square and two pairs of vertices, one on each side of the square, which are aligned to the two diagonals of the square. The smallest dissection is shown in Figure 2.5, it has six tetrahedra: In the top part

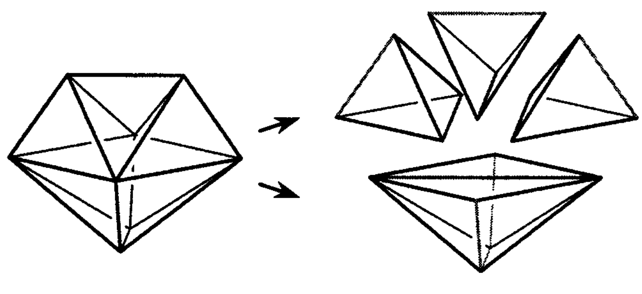

Figure 2.5: *A polytope whose minimal dissection is smaller than the minimal triangulation (dissection shown)*

there are the three tetrahedra spanned by the diagonal pointing away from us and the three edges on the top. Furthermore there are the three tetrahedra in the bottom part using the diagonal going across from us.

On the other hand each triangulation has at least seven tetrahedra: We have shown in Lemma 2.2 that a smaller triangulation (one with less than $8 + 2 - 3 = 7$ tetrahedra would have to use *both* diagonals of the square (as did the dissection). However, it is impossible to have both diagonals in a triangulation since they meet in a point which is a face to neither of them – the tetrahedra would not form a simplicial complex.

Hence the triangulation in Figure 2.6 with seven tetrahedra is minimal and larger than the minimal dissection.

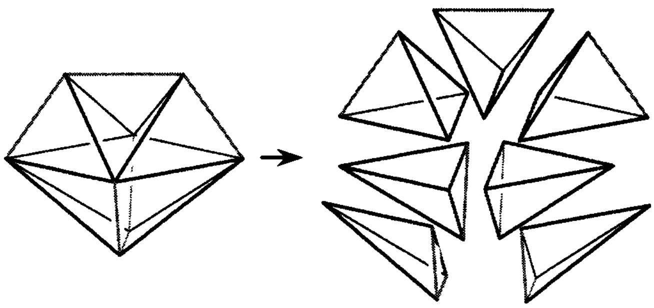

Figure 2.6: *A minimal triangulation of the polytope in Figure 2.5*

This construction can be generalized in the following fashion: We start with the square and build two vertex-edge chains of length m over its two diagonals, one going in the direction of each diagonal. We obtain P_m^{diss} : for instance in Figure 2.6 we were looking at P_2^{diss}. The $2m+2$ tetrahedra spanned by the edges of the vertex-edge chains and the suitable diagonal of the square make up a dissection of P_m^{diss}. (This is also the smallest dissection since each tetrahedron has two triangles of P_m^{diss} as facets, more cannot be achieved since no vertex of P_m^{diss} has degree 3.)

But since a triangulation cannot use both diagonals, by Lemma 2.2 it has to use at least $n + m - 3 = 3m + 1$ tetrahedra. This bound is tight: it can be achieved by a pulling triangulation: Choose all tetrahedra spanned by one of the four vertices of the inner square (say v) and a boundary triangle not incident to v. In Figure 2.6 v was chosen to be the vertex closest to the viewer. □

Notice that the dissection P_2^{diss} with 6 tetrahedra and 2 interior edges also proves that Lemma 2.3 cannot hold for dissections.

2.3.2 Face Lattice does not Suffice for Minimal Triangulation

The minimal size of a triangulation of a convex 3–polytope is *not* an invariant of the face lattice.

Theorem 2.5 *There is a simplicial convex 3-polytope with 10 vertices for which the minimal number of tetrahedra possible in a triangulation depends on its coordinates.*

PROOF: Consider the polytope $P = P_3^{diss}$ which is obtained by erecting two vertex-edge chains of length three over the two diagonals of a square (see the previous section). Now perturb the rightmost vertex by a slight downward movement and get the polytope P'. Figure 2.7 shows the minimal triangulation of P' (9 tetrahedra). But this is only possible because the four vertices of the tetrahedron in the middle are not coplanar. In the polyhedron P they are coplanar, and Lemma 2.2 tells us that P has a minimal triangulation of at least 10 tetrahedra. By coning from any of the six-valent vertices we can obtain such a triangulation. This proves the first part of Theorem 1.2. □

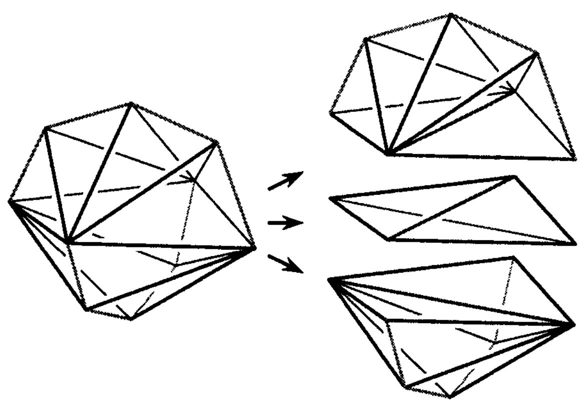

Figure 2.7: *The face lattice of this polytope does not determine the minimal triangulation*

This polytope is indeed the smallest example with this behavior. The proof can be found in Section 4 of [5].

Remark: In this case the degeneracy is *not* really necessary (as opposed to the last section):

First, it is clear that the polytope P' can be perturbed such that its vertices are in general position and such that its minimal triangulation is still 9.

Second, by pulling the rightmost vertex by a slight *upward* (followed by a perturbation which removes all degeneracies, but leaves all simplex orientations intact) we get a polytope P'' which also has a minimal triangulation of 10. Here is the reason: A triangulation using 9 tetrahedra would have to use both diagonals, but by Lemma 2.3 it could not use more than these two interior edges. However, the polytope has 16 triangles, so some tetrahedra in a size-9 triangulation must have more than one of these triangles on its boundary. There is no three-valent vertex, so there can only be two exterior triangles on each tetrahedron. If we look at all pairs of adjacent triangles and the edge shared by them we can identify two cases: Either the edge is on the vertex-edge chain and the tetrahedron is pierced by the "other" diagonal. This cannot happen in a triangulation. Or the edge is of the form aq_i and the opposite edge in the tetrahedron is an interior edge. But we had assumed that there are only two interior edges, the diagonals. So this is also impossible. Hence there is no triangulation with 9 tetrahedra and the smallest triangulation is the pulling triangulation from the last section with 10 tetrahedra.

2.3.3 Interior Points can make Minimal Triangulation Smaller

Theorem 2.6 *There is a family of convex polytopes P_m^{int} allowing a triangulation using an interior point which is smaller than every triangulation of P_m^{int} not using interior points. This gap is linear in the number of vertices.*

PROOF: Consider a triangular prism. We mark one diagonal edge of each quadrilateral face in a circular fashion (no two of the diagonals intersect) - see Figure 2.8. Over each quadrangular face we construct a vertex-edge chain of length m with the marked edge as diagonal. The vertices on the chains are placed sufficiently low over the base triangle of the prism such that the tetrahedron spanned by the top triangle and each of the points is pierced by the corresponding diagonal (this is closely related to the construction of the cupola from the Schönhardt polytope mentioned in the introduction which is made concrete in the next chapter). We obtain the polytope P_m^{int}.

The reader may verify that this is perfectly possible: First, errect a parabolic curve connecting the two outmost vertices of the new vertex-edge chain. Then identify a point on this parabola which lies below the intersection of this curve with the plane spanned by the diagonal and its "opposite" vertex of the top triangle. This point is low enough over the bottom triangle such that the the tetrahedron spanned by the point and the top triangle intersects the interior of the corresponding diagonal. Finally, place the points of the vertex-edge chain on the parabola between the constructed point and the bottom triangle.

All three vertex-edge chains allow us to invoke Lemma 2.2: So if a triangulation does not use any of the marked diagonals it will have at least $4m + 3$ tetrahedra (the polytope has $3m + 6$ vertices). On the other hand, we claim

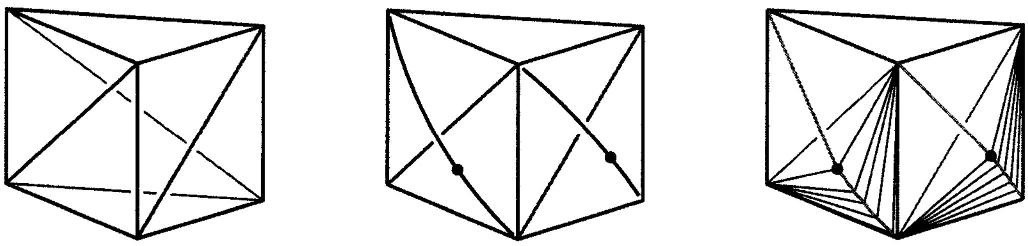

Figure 2.8: *The polytope P_m^{int} - a triangular prism patched with three vertex-edge chains*

that is impossible to have a triangulation which uses those three edges simultaneously. The reason is that the top face must belong to a tetrahedron, and if the fourth point is among the three vertices of the bottom triangle, the intersection with one of the diagonals is not a face to both of them. Finally, if the fourth point is a point of a vertex-edge chain, there is a bad intersection with the corresponding diagonal by construction. Hence, *every* triangulation must have at least $4m + 3$ tetrahedra.

For a small triangulation of P_m^{int} using an interior point triangulate the vertex-edge chains with the tetrahedra spanned by the diagonals and two consecute chain vertices. The remaining triangular prism now has to be triangulated using the marked diagonals. These diagonals split the quadrangular faces into two triangles - the bondary of the prism is triangulated. Choose any point in the interior of the prism and take all eight tetrahedra spanned by this point and one of the triangles in the boundary of the prism. Hence we have $3m + 11$ tetrahedra.

Putting it all together: For $m > 8$, our polytope P_m^{int} has a minimal triangulation which is larger than than the minimal triangulation using one interior point. The gap is $m - 8 = n/3 - 10$ which is linear in the number n of vertices. $\square$

Chapter 3

MINTRIANG(3) is Hard

3.1 Introduction

In the previous chapter we had glued vertex-edge chains to the three rectangular sides of a prism. Even more powerful effects can be observed if we perturb the vertices of the prism before the gluing. An appropriately perturbed version of the prism (the Schönhardt polytope) is another main ingredient (next to the vertex-edge chain). The patched version of this perturbed prism (the cupola) will help us to prove that many interior points might be needed for a minimal triangulation and that MINTRIANG is hard.

We introduce the non-convex *Schönhardt polytope* [9, 38, 39, 42, 51, 52]. It is a non-convex polytope obtained by twisting the top triangle of a triangular prism in a clockwise direction (see the first transition in Figure 3.1) The

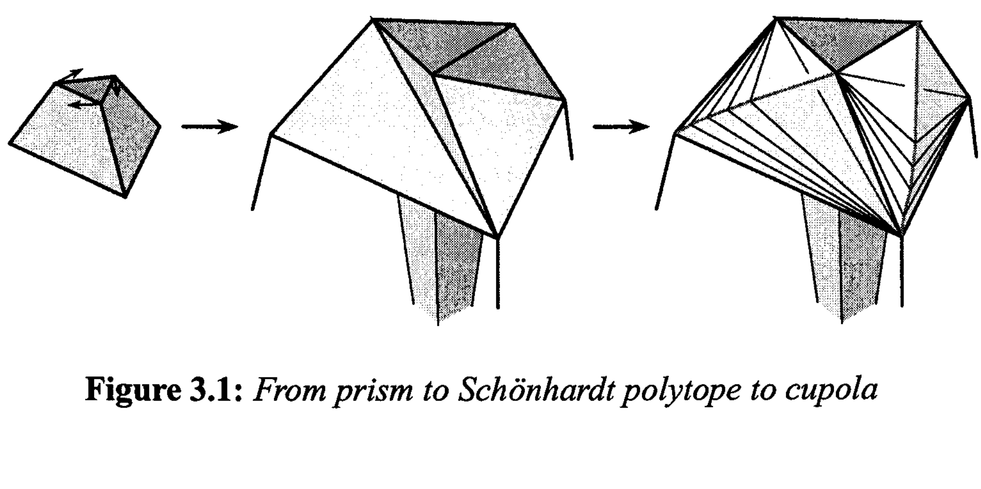

Figure 3.1: *From prism to Schönhardt polytope to cupola*

three quadrangular sides are then broken up and "bent in," thus creating the non-convex edges. The resulting polytope is non-convex and it cannot be triangulated: Much like in the last section, whichever point we would choose as fourth point of the tetrahedron having the top triangle as a face the tetrahedron would not be completely inside the polytope - the non-convex edges are "in the way."

Imagine the Schönhardt polytope part of a bigger polytope which stretches below its bottom face. This non-convex polytope can only be triangulated if the top triangle finds a fourth point with which it spans a tetrahedron which in turn is completely inside the polytope. This fourth point *has to lie* in the triangular cone defined by the faces adjacent to the top triangle as we will show in the next section. This cone is called *visibility cone.*

Now we convexify the Schönhardt polytope by attaching three circular vertex-edge chains opposite to the concavities. Thus, we create a convex polytope that satisfies the hypothesis of Lemma 2.2, and that we will call a *cupola* (see second transition in Figure 3.1). We can combine the properties of vertex-edge chains and of Schönhardt polytopes. Namely, in order to have a small triangulation of the whole polytope, the three diagonals of the Schönhardt polytope inside the cupola have to be used. But then, the tetrahedron containing the triangular top face of the Schönhardt polytope must not intersect the diagonals. The fourth vertex of this tetrahedron will have to lie in the triangular visibility cone we defined for the Schönhardt polytope. (This is only true if the vertices of the vertex-edge chain are placed sufficiently low over the bottom triangle; in this case the tetrahedra spanned by these vertices and the top triangle intersect the diagonals, thus cannot be used in a small triangulation.)

In summary, one can say that the convexified Schönhardt polytope – the cupola – as part of a bigger polytope forces small triangulations to contain a tetrahedron spanned by the triangular top face and a fourth vertex in a special cone determined by the Schönhardt polytope.

3.1.1 The Logical Polytope

In [51] Ruppert and Seidel used the Satisfiability problem (SAT, see p. 10) to prove that it is NP-complete to decide whether a non-convex polyhedron admits a triangulation. Their constructions used Schönhardt polytopes, and in particular their visibility cones, to do the transformation. In our case, because we need convexity and are talking about minimal (or at least small) triangulations, we glue cupolas, instead of Schönhardt polytopes. They are glued to a bigger polytope along their bottom faces. We call this polytope the *logical* *polytope* since for each logical (SAT) formula there will be a logical polytope which has a small triangulation if and only if the formula is satisfiable.

Similar to [51], we have variable cupolas and clause cupolas. The visibility cones of the variable cupolas contain only two *truth-setting* vertices, one for *false* and one for *true*. The visibility cones of the clause cupolas contain as many *literal* vertices as there are literals in the logical clause. Each variable must choose between a "true" or "false" value. Inside each clause at least one variable will be chosen to be true (to satisfy the clause). We model these logical choices by the geometric choices of which vertex in the visibility cone of a (variable/clause) cupola is used to triangulate the top face (see schematic Figure 3.2). In addition, our polytope satisfies some *blocking* conditions: the

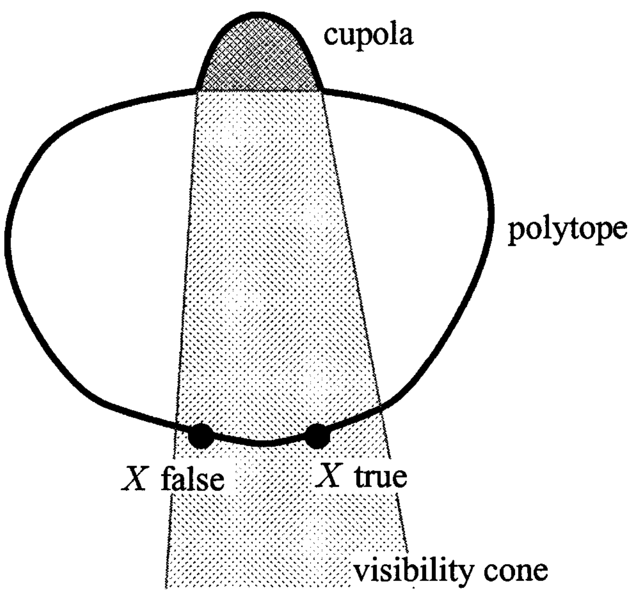

Figure 3.2: *One of the vertices is used to triangulate the cupola's top face, thus the choice between true and false*

tetrahedron spanned by the top face of a clause cupola and a literal vertex coming from a negated variable X_i will improperly intersect the tetrahedron spanned by the top face of the cupola of variable X_i and the truth-setting vertex corresponding to *true* (see Figure 3.3). In this way the choices made for the truth values of the variables and for the literals satisfying the clauses will be *consistent*.

In Section 3.3 we will give a detailed description of the logical polytope, of its building blocks and their interplay. We will see why there is no small triangulation if the logical formula is not satisfiable (relatively easy) and how

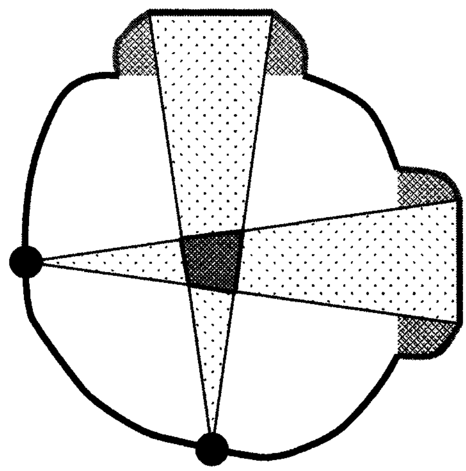

Figure 3.3: *Not both: Negated literal $\neg X_i$ triangulating its literal clause and true triangulating variable X_i*

to construct a small triangulation based on a satisfying truth assignment to the variables of the logical formula (harder).

Still missing at this point is the construction of the logical polytope given a logical formula. It is a priori not clear whether there is a polytope having the properties of the logical polytope. For the polynomial transformation (from SAT) however, we need to give an algorithm to compute the coordinates of the logical polytope. The binary encoding length of the polytope, as well as the running time of the algorithm, have to be polynomial in the encoding length of the SAT instance. Each step of the construction will be polynomial, this is a delicate point in the formal argument. We apply a sequence of these constructions (polynomially many). The coordinates of the vertices of the polytope are potentially singly exponential, but their binary encoding length is guaranteed to be polynomial. Section 3.4 is dedicated to this task.

3.1.2 Minimal Triangulations Needing Many Interior Points

We will now briefly describe how we can apply the technique of cupolas to showing that sometimes many interior points are necessary to get the minimal triangulation.

Consider a 3-polytope having many (say k) cupolas. Suppose that the visibility cone of each cupola contains no other vertices of the polytope except the vertices of the top triangle of the cupola and that visibility cones of different cupolas do not intersect. Assuming the vertex-edge chains of the cupola is very long (m a very big number), only by placing interior points in each visibility cone one can achieve a small triangulation. Hence the minimal triangulation of this polytope must use at least k interior points. In Section 3.5 we will make this argument precise.

3.2 Schönhardt Polytope and Cupola

3.2.1 The Technique of Gluing

We recall the notion of *beyond a face* (see [61] p. 78): A point p is *beyond* a face F of a polytope P if it (strictly) violates all inequalities defining facets of P containing F, but it strictly satisfies all other inequalities that define other facets of P. The polytope $P_{\text{beyond } F}$ is the (closure of the) set of all points beyond F (for a two-dimensional example see Figure 3.4). We denote by

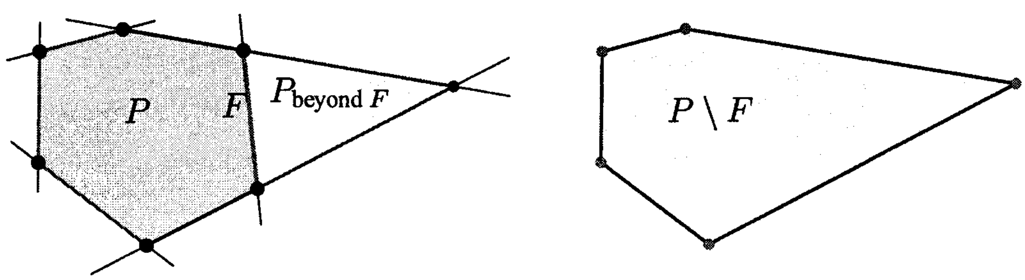

Figure 3.4: *The polytopes $P_{\text{beyond } F}$ and $P \setminus F$*

$P \setminus F$ the polyhedron defined by all facet–defining inequalities that do not hold with equality for all points in F. This is exactly $P \cup P_{\text{beyond } F}$. In our constructions we will often put one or more points beyond some face, and then take the convex hull. This will only destroy the facets containing this face, and introduce new ones containing the new points. We will say we *attach* (or *glue*) one polytope P to another polytope Q *along* facets F_P of P and F_Q of Q if $P \subseteq Q_{\text{beyond } F_Q}$ and $Q \subseteq P_{\text{beyond } F_P}$. It is important to observe that the result of the gluing, the convex hull of their union, contains both the face lattices of P and Q without, of course, F_P and F_Q (see Figure 3.5).

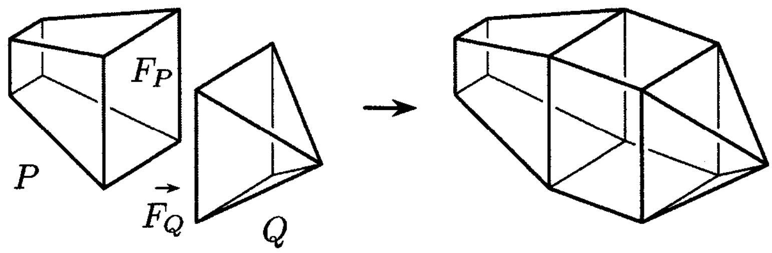

Figure 3.5: *The result of gluing P to Q is conv$(P \cup Q)$*

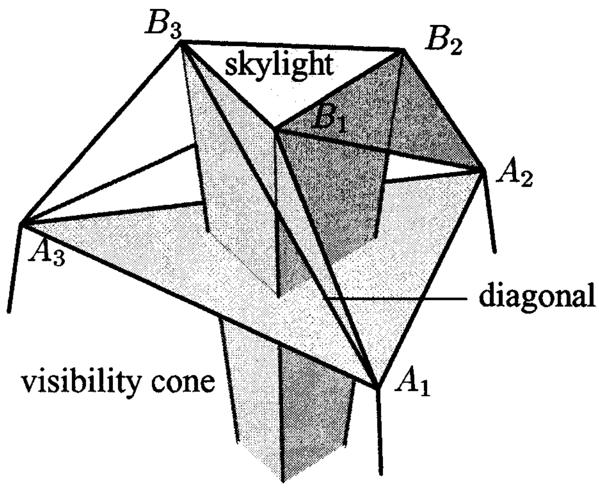

Figure 3.6: *Schönhardt polytope: its diagonals, skylight and visibility cone*

3.2.2 Schönhardt Polytope

The Schönhardt polytope is obtained from a triangular prism by twisting the triangular top face, breaking up and "bending in" the quadrangular faces (see the first transition in Figure 3.1). It was named after its first occurrence in [52] (see also [42]). For the notion of non-convex polytope and what it means to triangulate them we refer to [17].

Definition 3.1 *A* Schönhardt *polytope is a non-convex polytope with six vertices A_1, A_2, A_3, B_1, B_2, and B_3 and facets (A_1, A_2, A_3), (B_1, B_2, B_3), (A_1, B_1, A_2), (B_1, A_2, B_2), (A_2, B_2, A_3), (B_2, A_3, B_3), (A_3, B_3, B_1), and (B_3, B_1, A_1). At exactly the edges (B_1, A_2), (B_2, A_3), (B_3, A_1) the corresponding facets are to span an interior angle greater than π (the edges are* *said to be* reflex *or* non-convex*). These edges are called the* diagonals *of the Schönhardt polytope. The top face* (B_1, B_2, B_3) *is called the* skylight *of the Schönhardt polytope.*

Six points are said to be in Schönhardt position *if they are the vertices of a Schönhardt polytope. We say that the skylight is* visible *from a point* x *(or* x *is* able to see *the skylight, or* x *is a* viewpoint *of the skylight) if the tetrahedron spanned by* x *and the skylight does not intersect any of the diagonals in their relative interior. The* visibility cone *of the Schönhardt polytope is the triangular cone bounded by the planes* $B_1B_2A_2$, $B_2B_3A_3$, *and* $B_3B_1A_1$. *See Figure 3.6.*

The use of the word "skylight" is motivated by the idea that the skylight triangle is a glass window and light comes through it illuminating the interior of the Schönhardt polytope defining a cone of light. A point is visible by the skylight if none of the diagonals are "in the way." It is obvious that this non-convex polytope cannot be triangulated (without adding new points): The fourth point of the tetrahedron containing the skylight must be one of A_1, A_2, or A_3, but the diagonals "obstruct the view" of the skylight from these vertices. The visibility cone will be shown to contain all the visible points which are interesting in our construction.

3.2.3 Cupola

We will not use the Schönhardt polytope as is, but we will patch its sides with vertex-edge chains – and thereby convexify it – obtaining the *cupola.* Glued to a bigger polytope the cupola will force that certain tetrahedra occur in a small triangulations.

Definition 3.2 *A convex polytope* C *is called an* m-cupola *(or a* cupola *if the* m *is clear) if it has the following properties:*

1. *The vertices of* C *are* A_1, A_2, A_3, B_1, B_2, B_3, *and* q_k^i $(k = 0, \ldots, m+1, i = 1, 2, 3)$, *where the pairs* $q_0^i = A_i$ *and* $q_{m+1}^i = B_{i+1}$ *are identified (see Figure 3.7).*

2. *The vertices* A_1, A_2, A_3, B_1, B_2, *and* B_3 *are in Schönhardt position, and* (A_1, A_2, A_3) *(the bottom facet) and* (B_1, B_2, B_3) *(the skylight) are facets of* C.

3. *The other facets of* C *are* (B_i, q_k^i, q_{k+1}^i) *and* (A_j, q_k^i, q_{k+1}^i) *for* $k = 0, \ldots, m+1$, $i = 1, 2, 3$,

4. *the vertices q_k^i ($k = 1, \ldots, m$) lie on same the side of the plane $B_i A_{i+1} B_{i+2}$ as A_i ($i = 1, 2, 3$).*

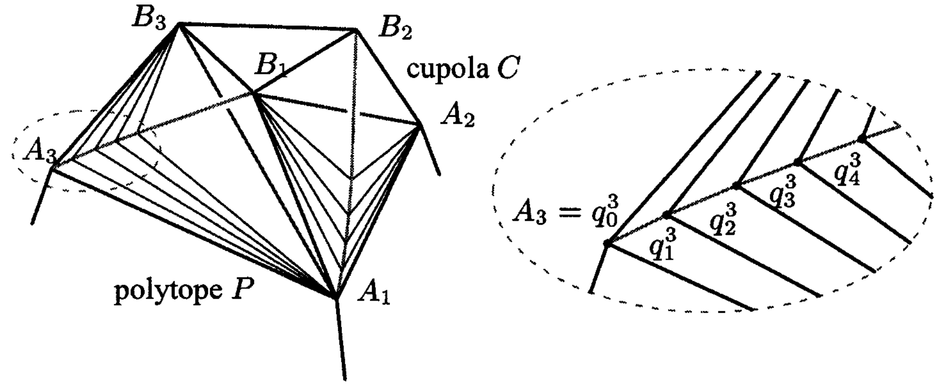

Figure 3.7: *A cupola as part of a larger convex polytope P*

Remark 3.3 *The combinatorial structure of the Schönhardt polytope and of the cupola are symmetric with a order of 3. The index arithmetic in the previous definition and in the sequel is therefore meant to be modulo 3. For instance, the diagonals are the edges (A_i, B_{i-1}) ($i = 1, 2, 3$).*

Let us digest this definition now. The basis of a cupola is a Schönhardt polytope. The sides of it are patched with vertex-edge chains such that (a) the resulting polytope is convex and (b) the Lemma 2.2 can be invoked forcing small triangulations to contain the diagonals as edges.

The last property of the cupola (Definition 3.2 (4)) means that the vertices of the vertex-edge chain are "low over" the bottom facet of the cupola. We need this condition for the following reason: Cupolas force small triangulations of the polytope to which they are glued to use the diagonals. From this we want to conclude that the skylight has to be triangulated by a vertex lying in the visibilty cone. If the vertices of the vertices of the vertex-edge chain q_k^1 are on the same side of plane $B_3 A_1 B_2$ as A_3 then the tetrahedron spanned by a q_k^3 and the skylight will be pierced by the diagonal $B_3 A_1$. The reader may check in Figure 3.7 that this feels right. Hence it cannot be one of these vertices triangulating the skylight. Item 2 of the next lemma makes this argument precise.

But if the skylight is triangulated by none of the vertices of the vertex-edge chains – and therefore by none of the vertices of the cupola – it has to be triangulated by a vertex lying beyond the bottom facet of the cupola. Item 3 of the next lemma states that in this case the triangulating vertex is in the visibility cone.

Lemma 3.4 ***Schönhardt Position.*** *Let A_1, A_2, A_3, B_1, B_2, B_3 be six points in Schönhardt position. We denote by $C_{A,B}$ the convex hull of the six points. Then*

1. *All orientations of simplices spanned by four of these six points are determined up to one global sign change. As a consequence, the six points are in convex position. Their convex hull $C_{A,B}$ is an octahedron that has (A_1, A_2, A_3) and (B_1, B_2, B_3) as facets, it has edges (A_i, B_{i+1}) ($i = 1, 2, 3$), and the line segments (A_i, B_{i-1}) ($i = 1, 2, 3$) are its diagonals*

2. *There are no points that can see the skylight (B_1, B_2, B_3) and, at the same time, (1) are beyond either of the edges (A_i, B_{i+1}) of $C_{A,B}$, and (2) are on the side of the plane $B_1A_2B_3$ opposite to B_2 or similarly for the analogous planes $B_1A_3B_2$, $B_2A_1B_3$ and the points B_3, B_1 respectively.*

3. *The visible points beyond the facet (A_1, A_2, A_3) of $C_{A,B}$ are exactly the points that are also in the visibility cone of the Schönhardt polytope.*

Take a look at Figure 3.8. It shows a point p outside of the visibility cone and thereby illustrates Item 3 of the previous lemma: a point below the face (A_1, A_2, A_3) is in the visibility cone if and only if the tetrahedron (p, B_1, B_2, B_3) pierces one of the diagonals.

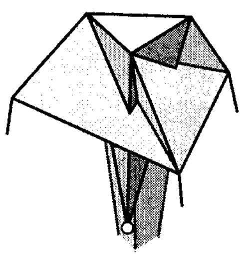

Figure 3.8: *A point outside the visibility cone*

Note also that Item 2 of the previous lemma is the reason for Item 4 of the definition of the cupola (Definition 3.2): The proof of this lemma is in Section 3.2.5.

3.2.4 How to Use Cupolas

We are now ready to give the statement about how we use cupolas:

Proposition 3.5 *Let C be an m-cupola which is part of a larger polytope P in the sense that the set $Q = closure(P - C)$ is a convex polytope and Q and C share the common facet (A_1, A_2, A_3). Let n be the number of vertices of P and n' be the number of vertices of Q.*

If T is a triangulation of P with the property that the fourth point of the tetrahedron containing the skylight of C is not *in the visibility cone of C, then there are at least $n + m - 3 = n' + 4m$ tetrahedra in the triangulation.*

PROOF.If the vertex triangulating the skylight of C is a vertex of on a vertex-edge chain of C, then it does not see the skylight by Definition 3.2.4 and Lemma 3.4.2. If it is in Q instead, then it has to be beyond the face (A_1, A_2, A_3) of C. Hence by Lemma 3.4.3 it cannot see the skylight either. Therefore the triangulation T does not use one of the diagonals. By Lemma 2.2 the number of tetrahedra is at least $n + m - 3$. Since by construction $n = n' + 3(m + 1)$, the number of tetrahedra is also at least $n' + 4m$. □

Proposition 3.5 stated that we get a large triangulation if we triangulate the skylight of a cupola by a vertex outside the visibility cone. Now we want to estimate how much smaller a triangulation is if we use a vertex v in the visibility cone instead. We give a relatively small triangulation of the cupola and of the space between the bottom face (A_1, A_2, A_3) of the cupola and the triangular face F of P with the help of the vertex v.

Proposition 3.6 *Let F be triangular face of a polytope P, and C an m-cupola attached to it according to Definition 3.2. Let v be a vertex of P in the visibility cone of C. Then there is a triangulation of conv($\{v\} \cup F \cup C$) with at most $3m + 16$ tetrahedra.*

PROOF.First of all, we triangulate along the vertex-edge chains using the tetrahedra $(B_i, A_{i+1}, q_k^i, q_{k+1}^i)$ for $i = 1, 2, 3$, and $k = 0, \ldots, m$.

After removing these tetrahedra, we are left with the union of the Schönhardt polytope on the vertices A_1, A_2, A_3, B_1, B_2, B_3, and the convex polytope conv($\{v\}, F, (A_1, A_2, A_3)$). This is a non-convex polytope with all edges, except the diagonals, being convex (easy conclusion from Lemma 3.4 and the construction). Since the specified vertex v is inside the visibility cone, it sees all facets of this polytope, except the three facets it is incident to, from the *interior*. In particular, we can form tetrahedra of all these facets and v

and none of them intersect badly. They are 7 tetrahedra for the facets of the Schönhardt polytope (since we do not count the bottom face) and at most 6 for the rest (the convex hull of F and (A_1, A_2, A_3) has—by a planar graph argument—at most $2 \cdot 6 - 4 = 8$ facets, subtracting 2 for F and (A_1, A_2, A_3) gives 6). □

It is this $3m$ in contrast to the $4m$ in Proposition 3.5 which makes this way of triangulating optimal for large m. We give more details on the use of these propositions in Section 3.5 when we look at an actual logical polytope with many cupolas.

3.2.5 Proof of the Schönhardt Position Lemma

In the proof of Lemma 3.4 we have to make argument about the relative position of points with respect to planes spanned by points. For instance it is necessary (but not sufficient) for a line segment to pierce a triangle that the endpoints of the segment lie on opposite sides of the triangle. This kind of information is captured by the *oriented matroid* of the points. In fact, in our proof we only use the oriented matroid properties of the relative positions of the points.

For the theory of oriented matroids we refer to [12] and [61]. Here we only sketch the necessary definitions and how they are related to the notion of visibility. The orientation of a simplex (x_1, x_2, x_3, x_4), is defined as

$$[x_1, x_2, x_3, x_4] = \text{sign det} \begin{pmatrix} x_1 & x_2 & x_3 & x_4 \\ 1 & 1 & 1 & 1 \end{pmatrix}.$$

It is easy to see that this definition of orientation reflects our intuition of orientation perfectly. All such orientations make up the *chirotope* of an oriented matroid (see page 123 in [12]).

Given the oriented matroid of points $x_1, \ldots, x_n$ in d–space, its circuits are functions $C : \{x_1, \ldots, x_n\} \mapsto \{+, -, 0\}$ that correspond to so–called minimal Radon partitions. This means that the convex hulls of $C^+ = \{x_i | C(x_i) = +\}$ and $C^- = \{x_i | C(x_i) = -\}$ intersect in their relative interiors, and C^+ and C^- are minimal at that. It is easy to check that the function

$$C(x) = \begin{cases} (-1)^i \cdot [\overbrace{x_1, \ldots, x_{d+1}}^{\text{omit } x_i}] & \text{, if } x \in \{x_1, \ldots, x_{d+1}\}, \\ 0 & \text{otherwise,} \end{cases}$$

defines a circuit if it is not identical 0. In fact, all circuits can be obtained this way. We will compute circuits to use an argument of the following form: x does not see the skylight if and only if there is a circuit such that the positive part is one of the diagonals and negative part is the set containing x and a subset of vertices of the skylight. Since then the tetrahedron spanned by x and the skylight is pierced by the diagonal.

Important tools to compute simplex orientations are the Grassmann–Plücker relations (see Section 2.4 in [12]): For points a, b, $x_1, \ldots, x_4$ they state that the set of signs

$$\begin{array}{ll} \{ & [a,b,x_1,x_2] \cdot [a,b,x_3,x_4], \\ & -[a,b,x_1,x_3] \cdot [a,b,x_2,x_4], \\ & [a,b,x_1,x_4] \cdot [a,b,x_2,x_3] \quad \} \end{array}$$

is either identical 0 or contains both a $+$ and a $-$. The Grassmann–Plücker relations follow easily from the corresponding determinant equation (try it!). The typical use of the Grassmann–Plücker relations is to deduce one orientation when the others are known. We can read the orientations of some of the different tetrahedra from two-dimensional projections (drawings) of the point configurations as in Figure 3.6. We use a left–hand–rule coordinate system, i.e., we decide whether the triangle (x_1, x_2, x_3) is oriented counterclockwise $(+)$ or not $(-)$, also if x_4 is on our side of the plane spanned by x_1, x_2, and x_3 $(+)$ or not $(-)$, and multiply these two signs to obtain the orientation $[x_1, x_2, x_3, x_4]$.

PROOF OF LEMMA 3.4.

1. In a Schönhardt polytope, the simplices (A_1, A_2, A_3, B_1), and (A_1, A_2, A_3, B_2) have the same orientation since edges (A_1, A_2) and (A_2, A_3) are both incident to facet (A_1, A_2, A_3) and they are both non-reflex edges.

By the this argument, going around the boundary of a Schönhardt polytope, keeping in mind which edges are reflex, we can determine the orientation of 12 simplices up to one global sign change (there are 12 edges). But there are $\binom{6}{4} = 15$ simplices formed by the vertices of the Schönhardt polytope. The remaining three simplices are (A_1, A_2, B_2, B_3), (A_2, A_3, B_1, B_3), (A_1, A_3, B_1, B_2). The signs are determined by the following Grassmann-Plücker relations: For (A_1, A_2, B_2, B_3) take $a = A_1$, $b = A_2$, $x_1 = A_3$, $x_2 = B_1$, $x_3 = B_2$,

$x_4 = B_3$ (the other two by circular index shift). Then:

$$\begin{aligned}\{ \quad & \overbrace{[A_1,A_2,A_3,B_1]}^{-}\cdot\overbrace{[A_1,A_2,B_2,B_3]}^{?},\\ & -\overbrace{[A_1,A_2,A_3,B_2]}^{-}\cdot\overbrace{[A_1,A_2,B_1,B_3]}^{+},\\ & \overbrace{[A_1,A_2,A_3,B_3]}^{-}\cdot\overbrace{[A_1,A_2,B_1,B_2]}^{-} \quad \} \supseteq \{+,-\}.\end{aligned}$$

This forces $[A_1,A_2,B_2,B_3] = +$.

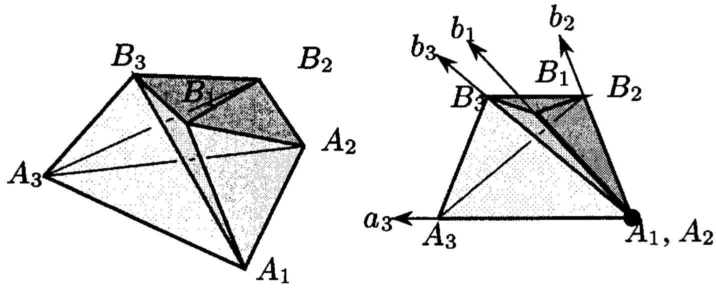

Figure 3.9: *Two views of the Schönhardt polytope, the second view is along edge* (A_1A_2)

There is a nice way that the use this Grassmann-Plücker relation can be pictured. Namely, an equivalent formulation is that there are rays a_3, b_1, b_2, and b_3 emitting from the origin of the plane such see in the direction of x ray y is to the left if and only if the sign $[A_1,A_2,X,Y] = +$. (The rays come from a special view of the configuration along the edge (A_1,A_2). This corresponds to the fact that the signs involving A_1 and A_2 capture the projection of the point configuration along the line (A_1,A_2).) But the signs say that a_3 has all b_i right from it, but b_3 is to the left of b_1 and b_2 is to the right of b_1, hence b_3 must be to the left of b_2. Therefore $[A_1,A_2,B_2,B_3] = +$.

From the chirotope information it is easy to check that all vertices are in convex position (see description of how to read the facets of the convex hull from the chirotope in Chapter 3 of [12]), and that their convex hull $C_{A,B}$ is indeed an octahedron.

2. We will show that if a point x lies beyond A_1B_2 of $C_{A,B}$, on the side of $B_1A_2B_3$ opposite to B_2 , then (B_1,A_2) and the triangle (B_2,B_3,x) form a minimal Radon partition in the set of vertices A_1, A_2, A_3, B_1, B_2, B_3, and x,

hence have an interior point in common. This means x cannot see the skylight. For this, we compute the following orientations:

$$\begin{array}{lcll}
-[B_1,B_2,B_3,x] & = & +, & \text{since } (B_1,B_2,B_3) \text{ is a facet of } C_{A,B} \setminus (A_1,B_2), \\
+[A_2,B_2,B_3,x] & = & +, & \text{since } (A_2,B_2,B_3) \text{ is a facet of } C_{A,B} \setminus (A_1,B_2), \\
-[A_2,B_1,B_3,x] & = & -, & \text{from the assumption on } x, \\
+[A_2,B_1,B_2,x] & = & -, & \text{from the Grassmann–Pücker relation below}, \\
-[A_2,B_1,B_2,B_3] & = & -, & \text{from Part (1).}
\end{array}$$

The necessary Grassmann–Plücker relation is the one with $a = B_1$, $b = B_2$, $x_1 = B_3$, $x_2 = A_1$, $x_3 = A_2$, and $x_4 = x$ such that

$$\begin{array}{ll}
\{ & \overbrace{[B_1,B_2,B_3,A_1]}^{-} \cdot \overbrace{[B_1,B_2,A_2,x]}^{?}, \\
 & -\overbrace{[B_1,B_2,B_3,A_2]}^{-} \cdot \overbrace{[B_1,B_2,A_1,x]}^{-}, \\
 & \overbrace{[B_1,B_2,B_3,x]}^{-} \cdot \overbrace{[B_1,B_2,A_1,A_2]}^{+} \quad \} \supseteq \{+,-\}
\end{array}$$

forces $[B_1,B_2,A_2,x] = -$.

3. If x is in the visibility cone V, then it is, by part (2) of this lemma, on the same side as B_3 with respect to the plane $B_1A_2B_2$. Hence A_2 is on opposite side of B_3 with respect to the plane B_1B_2x. Therefore, the relative interior of the convex hull of B_1 and A_2 lies strictly on one side of the plane B_1B_2x, and the tetrahedron (B_1,B_2,B_3,x) on the other side of this plane. Therefore those two point sets cannot have points in common. By symmetry it follows that the other two diagonals do not obstruct any point of V from seeing the skylight either.

Assume now that a point x is beyond face (A_1,A_2,A_3), but outside of V, i.e., for instance on the A_1 side of the plane $B_1B_2A_2$. We claim that the pair $\{B_1,A_2\}$, $\{B_2,B_3,x\}$ forms a circuit in the oriented matroid of the point configuration of the vertices of $C_{A,B}$ and x. This means that the triangle (B_2,B_3,x) is pierced by the diagonal (B_1,A_2) in the relative interior, hence x is not visible.

$$\begin{array}{lcll}
-[B_2,B_3,A_2,x] & = & -, & \text{since } (B_2,B_3,A_1) \text{ is a facet of } C_{A,B} \setminus (A_1,A_2,A_3), \\
+[B_1,B_3,A_2,x] & = & +, & \text{from the Grassmann–Plücker relations below}, \\
-[B_1,B_2,A_2,x] & = & +, & \text{from the assumption on } x, \\
+[B_1,B_2,B_3,x] & = & -, & \text{since } (B_1,B_2,B_3) \text{ is a facet of } C_{A,B} \setminus (A_1,A_2,A_3), \\
-[B_1,B_2,B_3,A_2] & = & +, & \text{from Part (1).}
\end{array}$$

In this case, we have to apply the Grassmann–Plücker relations twice to get $[B_1, B_3, A_2, x] = +$. First we deduce $[A_1, A_2, B_3, x] = -$ from the Grassmann–Plücker relation with $a = A_1, b = A_2, x_1 = A_3, x_2 = x, x_3 = B_2, x_4 = B_3$:

$$\begin{aligned} \{ \quad & \overbrace{[A_1, A_2, A_3, x]}^{-} \cdot \overbrace{[A_1, A_2, B_2, B_3]}^{-}, \\ - & \overbrace{[A_1, A_2, A_3, B_2]}^{+} \cdot \overbrace{[A_1, A_2, x, B_3]}^{?}, \\ & \overbrace{[A_1, A_2, A_3, B_3]}^{+} \cdot \overbrace{[A_1, A_2, x, B_2]}^{+} \quad \} \supseteq \{+, -\} \end{aligned}$$

Now we use this orientation to formulate $a = A_2, b = B_3, x_1 = A_1, x_2 = B_1$, $x_3 = B_2, x_4 = x$:

$$\begin{aligned} \{ \quad & \overbrace{[A_2, B_3, A_1, B_1]}^{+} \cdot \overbrace{[A_2, A_3, B_2, x]}^{-}, \\ - & \overbrace{[A_2, B_3, A_1, B_2]}^{+} \cdot \overbrace{[A_2, B_3, B_1, x]}^{?}, \\ & \overbrace{[A_2, B_3, A_1, x]}^{-} \cdot \overbrace{[A_2, B_3, B_1, B_2]}^{+} \quad \} \supseteq \{+, -\} \end{aligned}$$

in order to get the desired $[A_2, B_3, B_1, x] = -$. □

3.3 The Transformation from SAT

It is our intention to model the well-known satisfiability problem (SAT) using the visibility cones of cupola polytopes. Just as Ruppert and Seidel did in [51], from now on we will restrict our attention to *special SAT instances: each variable appears exactly three times, twice unnegated and once negated.* This is not really necessary, but simplifies explanations. We will argue in Section 3.3.1 why this restriction is valid. The formula

$$f = (X_1 \vee \neg X_2 \vee X_3 \vee \neg X_4) \wedge (\neg X_1 \vee X_2 \vee \neg X_3 \vee X_4) \wedge (X_1 \vee X_2 \vee X_3 \vee X_4)$$

is such a special SAT formula. The figures in this section will correspond to this particular instance.

In Section 3.3.1, we will introduce the formal definition of the family of *logical polytopes* $\mathcal{P}_{f,m}$ associated to a given logical formula f and a natural number m. This number m will denote the length of the vertex-edge chains of the cupolas we use. In order to define logical polytopes, we will specify the face lattice of a *frame polytope*, reveal where the m-cupolas are glued, and then give more conditions on the positions of the vertices relative to each other. To achieve our goal we need two lemmas. The first lemma ensures that logical polytopes exist and that we can construct them in polynomial time (see the proof in Section 3.4.4). The second lemma assures that, among all logical polytopes of a fixed logical formula, the size m of the vertex chains in cupolas can be adjusted to (1) be polynomial in the number of logical variables and clauses, and (2) to be large enough to guarantee the equivalence between logical satisfiability of the logical formula and small triangulations of the logical polytope. We will prove this second lemma in Section 3.3.2.

Lemma 3.7 *There is a polynomial algorithm that, given any positive integer m and a logical formula f containing C clauses and V variables, produces a logical polytope $P \in \mathcal{P}_{f,m}$ with m vertices on each vertex-edge chain. The number of vertices of P is bounded by a polynomial in m, C, and V. Also, the coordinates of the vertices have binary encoding length polynomial in m, C, and V.*

Lemma 3.8 *Let f be a logical formula containing C clauses and V variables. There exists a polynomial $m(C, V)$ with integer coefficients such that for $m = m(C, V)$ and for any logical polytope $P \in \mathcal{P}_{f,m}$ the following is true: P admits a triangulation with $\leq K = n + m - 4$ tetrahedra if and only if there is a satisfying truth assignment to the variables of the formula f.*

Finally, using these two properties, we are ready for the proof of the main result.

PROOF OF THEOREM 1.5: The problem is clearly in NP: checking whether a collection of tetrahedra is indeed a triangulation of the polytope P needs only a polynomial number of calculations. Every pair of tetrahedra is checked for proper intersection (in a common face or not at all), and the sum of the volumes equals the volume of P (computable for instance by a pulling triangulation of the polytope, see Section 2.3.1). Also the size of triangulations of a given polytope is bounded by a polynomial in n of degree two (this follows from the well-known upper bound theorem, for details see [49]).

By Lemma 3.7, from a given logical formula f on V logical variables and C clauses, we can construct a logical polytope $P \in \mathcal{P}_{f,m(C,V)}$ of encoding length polynomial in V and C. Hence, by Lemma 3.8 there is a polynomial transformation that establishes the polynomial equivalence of a solution for the SAT problem and the existence of small triangulations of P (small means of size less than K as given in the statement). This completes the proof. □

3.3.1 Logical Polytopes

First of all, we want to argue why the restriction to the *special SAT formulas* (every variable appears exactly twice unnegated and once negated) is NP-complete: The SAT problem remains NP-complete even for instances where each variable or its negation appear at most three times (see references on page 259 in [31]). In addition, a change of variables can be used to change a non-negated variable into a negated variable if necessary. Also note that if a variable appears only negated or only positive the variable and the clauses that contain it can be discarded. Finally, if a variable appears exactly once positive and exactly once negated then it can be eliminated by combining the two clauses that contain the two variables into one.

Now, we want to define the family of logical polytopes $\mathcal{P}_{f,m}$ for a given logical formula f and a given positive integer number m. We start by describing its face lattice. To prevent a possible confusion we remark that our vertices will be labeled by the letters c_j when the are related to logical clauses, x_j when they are related to logical variables, and z_j when the vertex is auxiliary. Points always have subscripts and/or subscripts thus should not be confused with their coordinate-entries (x, y, z).

In a logical polytope there will be an m-cupola for each clause and one for each variable and its negation. The cupolas will be glued to a *frame polytope* which resembles a wedge. Look carefully at Figure 3.10 for an example of the overall structure.

Figure 3.11 gives a view of the lower hull of the frame polytope. The sharp part of the wedge consists of $2C + 1$ vertices (where C is the number of clauses) $c_0, \ldots, c_{2C}$. We call this part of the frame polytope the *spine*. We attach the *clause cupola* associated with clause i to the triangle $(c_{2i}, c_{2i+1}, c_{2i+2})$ (shaded in the picture).

On top of this wedge structure we will put a series of *roofs*. They are triangular prisms, spanned by the two triangles (z^i_T, z^i_F, z^i_A) and (z^i_L, z^i_R, z^i_B), one for every variable X_i of the logical formula. The *variable cupolas* will be attached to the triangular facet (z^i_L, z^i_R, z^i_B), the *back gables* (the triangular faces are shaded in Figure 3.12).

The variable cupola of variable X_i is such that its visibility cone contains exactly the front vertices vertices z^i_T and z^i_F. We will use these cupolas to

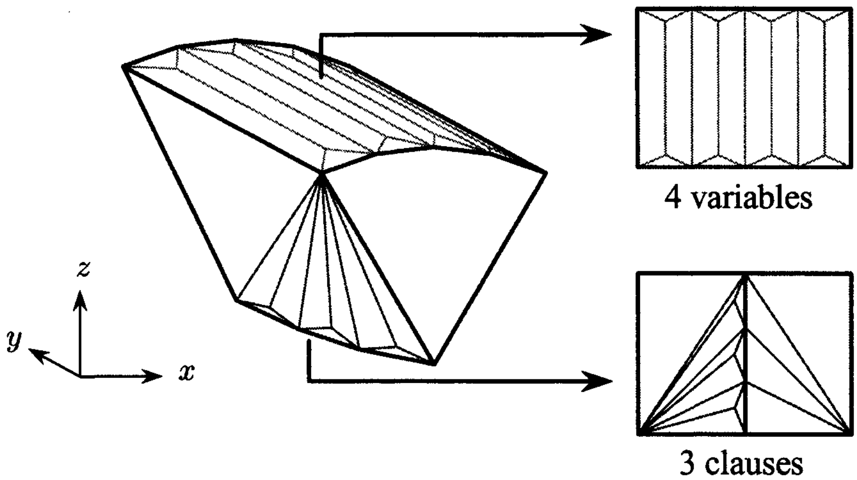

Figure 3.10: *Sketch of a logical polytope*

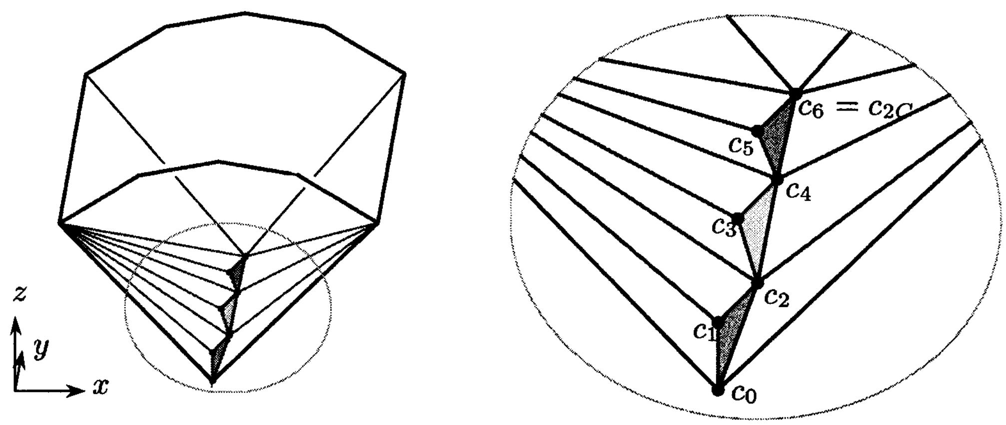

Figure 3.11: *The spine of the wedge: here the clause cupolas are attached*

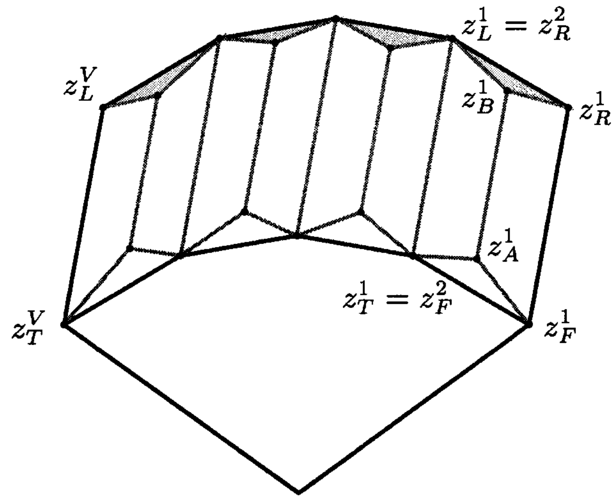

Figure 3.12: *The roofs, back gables shaded*

read from a small triangulation of the polytope the logical value of variables according with the following rule: if the truth-setting vertex z_T^i associated to the ith logical variable is used to triangulate the skylight of the cupola for variable i, then we set $X_i = true$. If the truth-setting vertex used to triangulate the skylight of the cupola for variable i is instead z_F^i then $X_i = false$.

Beyond the quadrilateral face containing z_T^i we will place the *literal vertices* x_1^i and x_2^i which corresponds to the positive occurrences of X_i in the logical formula. Beyond the other quadrilateral face we will place the other literal vertex $\overline{x_3^i}$ which correspond to the negated occurrence of X_i. These vertices are in the visibility cones of the three cupolas of the clause where variable X_i or its negation appears.

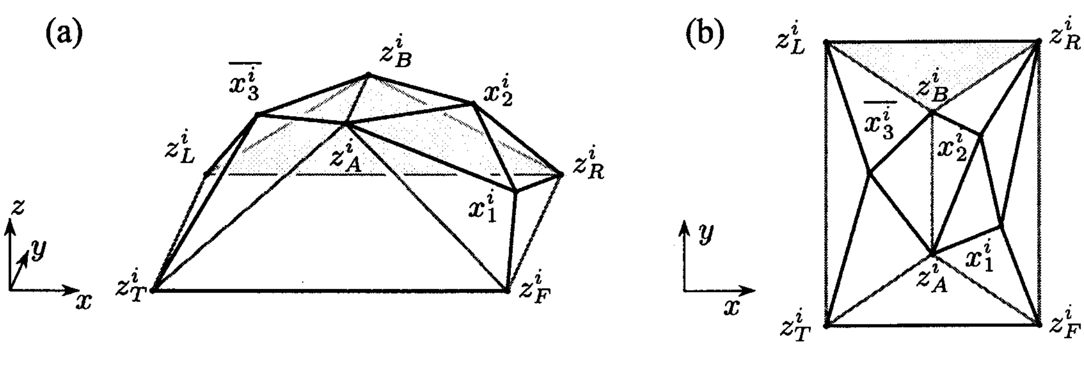

Figure 3.13: *A roof, back gable shaded, z-coordinate superelevated*

We list the five conditions on logical polytopes which are necessary for the transformation to work in both ways, i.e. a small triangulation yields a satisfying truth assignment for our logical formula and vice versa.

Definition 3.9 *For a given logical formula, the family $\mathcal{P}_{f,m}$ of* logical polytopes *is the set of all three-dimensional polytopes P that satisfy the following conditions:*

1. (Convexity and Face Lattice) *The logical polytope must be convex, and the face lattice is as we just described it. In particular, several m-cupolas are part of the polytope, one for each clauses and variable in f.*

2. (Visibility) *The literal vertices x_1^i, x_2^i, and $\overline{x_3^i}$ are vertices in the visibility cone associated to their respective clause m-cupolas, but of no other clause visibility cone. The vertices z_T^i,z_F^i are the only vertices in the visibility cones of the ith variable m-cupola.*

3. (Blocking) *This constraint ensures that the assignment of true or false values for variables is done consistently, i.e. the positive (negative) literals can be used to make their clauses true if and only if the variable is set true (false).*

   Concretely, the tetrahedron spanned by z_F^i and the skylight of the m-cupola of variable X_i intersects the interior of the tetrahedron spanned by x_1^i (by x_2^i) and the skylight of the clause m-cupola corresponding to x_1^i (to x_2^i). Also the tetrahedron spanned by z_T^i and the skylight of the m-cupola of variable X_i intersects the interior of the tetrahedron spanned by $\overline{x_3^i}$ and the skylight of the clause m-cupola corresponding to it. See Figure 3.14 for an example.

4. (Non-blocking) *Using the vertex z_T^i to triangulate the interior of the i-th variable m-cupola should not prevent the non-negated literal vertices from seeing their associated m-cupolas. Concretely, if j is the clause corresponding to the literal vertex x_1^i, then tetrahedra $(z_T^i, z_L^i, z_R^i, z_B^i)$ and $(x_1^i, c_{2j-2}, c_{2j-1}, c_{2j})$ do not intersect at all. The canonical analogue shall hold for x_2^i and $\overline{x_3^i}$ (for $\overline{x_3^i}$ replace z_T by z_F).*

5. (Sweeping) *Because we intend to follow the same triangulation procedure which was proposed by Ruppert and Seidel [51], and which we will explain in Section 3.3.2, we will need that*

   (a) the variable x_1^i is to the "left" (negative x direction) of the planes $c_{2k-1}c_{2k}z_F^i$, $c_{2k}c_{2k+1}z_F^i$, and $c_{2k-1}c_{2k+1}z_F^i$ for $0 \leq k \leq C-1$.

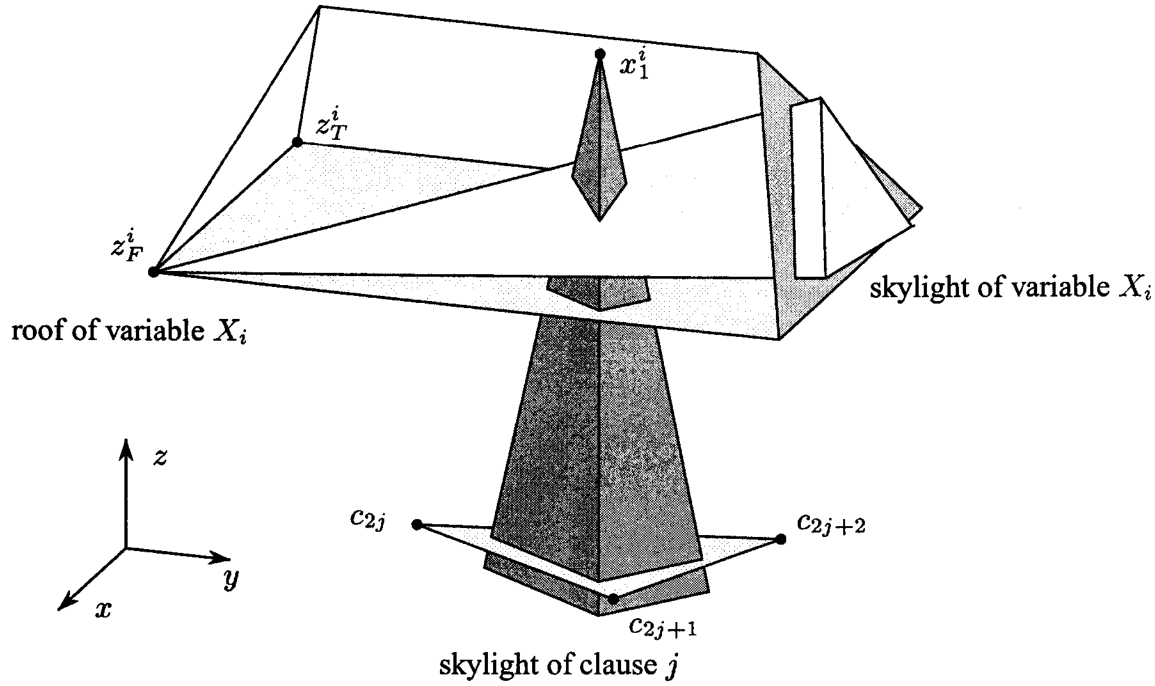

Figure 3.14: *Blocking for consistent logical values*

(b) x_2^i *is to the "left" of the planes* $c_{2k-1}c_{2k}x_1^i$, $c_{2k}c_{2k+1}x_1^i$, *and* $c_{2k-1}c_{2k+1}x_1^i$ *for* $0 \leq k \leq C-1$.

(c) $\overline{x_3^i}$ *is to the "left" of the planes* $c_{2k-1}c_{2k}z_F^i$, $c_{2k}c_{2k+1}z_F^i$, *and* $c_{2k-1}c_{2k+1}z_F^i$ *for* $0 \leq k \leq C-1$.

(d) z_T^i *is to the "left" of the planes* $c_{2k-1}c_{2k}x_2^i$, $c_{2k}c_{2k+1}x_2^i$, $c_{2k-1}c_{2k+1}x_2^i$, $c_{2k-1}c_{2k}\overline{x_3^i}$, $c_{2k}c_{2k+1}\overline{x_3^i}$, *and* $c_{2k-1}c_{2k+1}\overline{x_3^i}$ *for* $0 \leq k \leq C-1$.

3.3.2 Using the Logical Polytope

Proof of Lemma 3.8: If a triangulation T of the polytope has $\leq n+m-4$ tetrahedra, then by Proposition 3.5 the skylight of each cupola is triangulated by a vertex in the visibility cone of the cupola. In particular, one of z_F^i and z_T^i is chosen to triangulate the cupola corresponding to variable X_i for each i. We claim that assigning to X_i the truth value according to this choice ($z_F \mapsto false$, $z_T \mapsto true$) satisfies all clauses of the formula.

Each clause cupola skylight is triangulated by one of the literal vertices, say clause j by the positive literal vertex x_1^i (or x_2^i). By the blocking conditions, it cannot be the case that the variable skylight of X_i is triangulated by z_F^i since these tetrahedra would intersect badly. So we had set X_i to $true$. Having x_1^i

(or x_2^i) in clause j's visibility cone meant that variable X_i appears unnegated in this clause. If the skylight of clause cupola j is triangulated by $\overline{x_3^i}$, by the same argument X_i was set to false, and clause j satisfied by the literal $\neg X_i$. Hence all clauses are satisfied.

Now we need to prove the converse. If there is a $true$-$false$ assignment that satisfies all logical clauses we must find a triangulation that has no more than K tetrahedra. For that we construct a "small" triangulation. There are three different kinds of tetrahedra: the ones triangulating the cupolas, the ones triangulating the roofs, and the ones triangulating of the rest of the wedge. We know how to triangulate a cupola if we know a vertex in its visibility cone (see the proof of Proposition. 3.6). For the rest we will now follow a *sweeping procedure* which was first described by Ruppert and Seidel [51].

The sweeping triangulation proceeds by triangulating "slices" that correspond to the different variables X_1 to X_V, i.e. from right to left. The variable roofs are arranged sequentially for exactly this purpose. A slice is roughly speaking the part of the tetrahedra between a roof and vertices of the spine. After the ith step of the process the partial triangulation will have triangulated the first i slices. The part of the boundary of the partial triangulation that is inside the logical polytope will form a triangulated disk. We will call it the *interface* following the convention of Ruppert and Seidel. It contains the following triangles:

$$(z_T^i, c_{2C}, z_L^i) \text{ and } \begin{cases} (z_T^i, c_{2j-2}, c_{2j}) & : \text{ if clause } j \text{ is satisfied by one of the first } i \text{ variables, or} \\ (z_T^i, c_{2j-2}, c_{2j-1}) \\ \text{and } (z_T^i, c_{2j-1}, c_{2j}) & : \text{ otherwise,} \end{cases}$$

for all $j = 1, \ldots, C$. Before the first step, the partial triangulation is empty. After the last step the partial triangulation will cover the whole logical polytope. In general, the vertices of the ith roof will see all triangles of the interface and will be used as apexes to form new tetrahedra to add to the current partial triangulation. This way the interface will slowly move from right to left.

Now we describe in detail the triangulation step for the ith variable X_i. Since we are only concerned with roof vertices in roof i, we will drop all superscripts. The triangulation step depends on whether X_i is set $true$ or $false$ in the satisfying assignment. Let us consider first the case $X_i = true$:

The point z_T is used to triangulate the interior of the variable cupola associated to X_i according to Proposition 3.6. From z_T we also form six tetrahedra with the following triangles: $(z_L, \overline{x_3}, z_B)$, $(\overline{x_3}, z_B, z_A)$, (z_B, z_A, x_2), (z_B, x_2, z_R), (z_A, x_1, x_2), and (x_1, z_A, z_F).

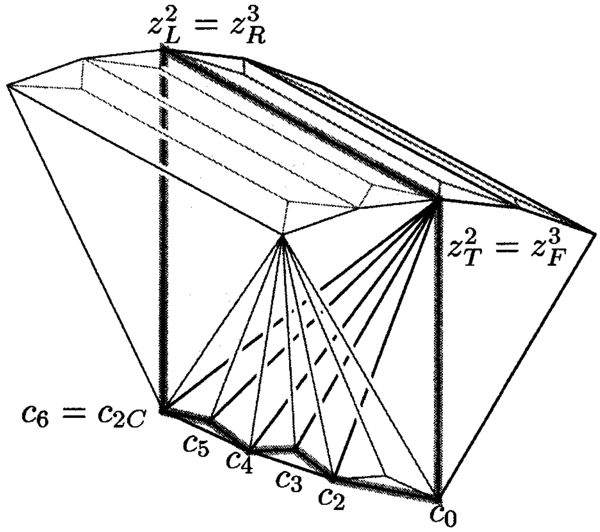

Figure 3.15: *The interface after step 2*

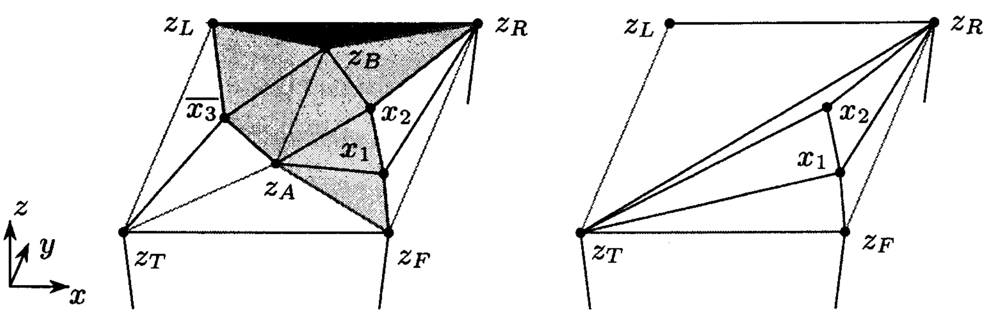

Figure 3.16: *Removing the tetrahedra spanned by* z_T *and the shaded triangles*

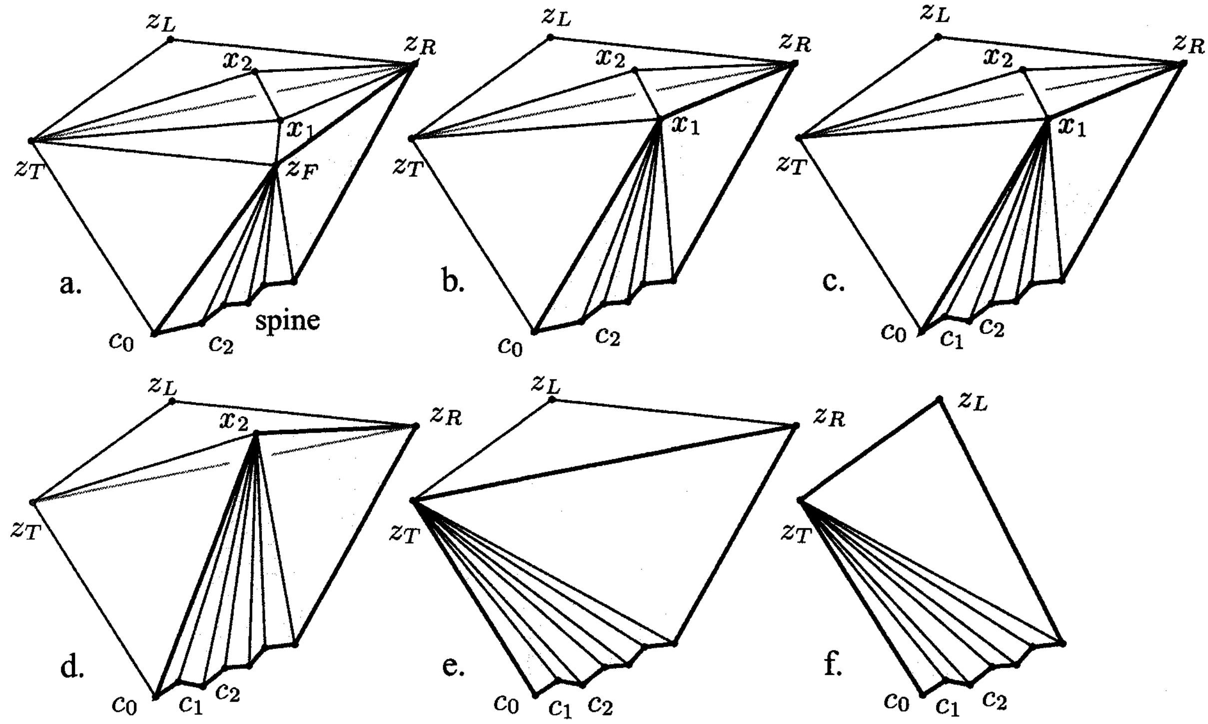

Figure 3.17: *The sweep*

Now we come to the part of the triangulation which gave the sweeping procedure its name. We form the tetrahedra between x_1 and the current interface triangles. This is possible by part (a) of Condition 5. We also use the tetrahedron (x_1, z_T, c_0, z_F). This is illustrated in the transition from a. to b. in Figure 3.17. The interface advances to x_1, i.e. if (z_F, c_j, c_k) was an interface triangle before, now (x_1, c_j, c_k) is an interface triangle. Also (z_F, c_{2C}, z_R) is replaced by the triangle (x_1, c_{2C}, z_R).

Since X_i is set to *true* we can use x_1 to triangulate its clause cupola according to Proposition 3.6. We only do this if the clause cupola has not been previously triangulated using an other literal vertex. Condition 2 ensures that x_1 is in the visibility cone of the clause cupola coming from the clause that contains the unnegated literal X_i. Furthermore, Condition 4 ensures that we can actually perform this triangulation of the clause cupola without badly intersecting the tetrahedra of the variable cupola. In Figure 3.17.c. we see that if x_1 is used to triangule clause j's cupola, then the interface triangle (x_1, c_{2j-2}, c_{2j}) is replaced by the two triangles $(x_1, c_{2j-2}, c_{2j-1})$ and (x_1, c_{2j-1}, c_{2j}).

We repeat this procedure with x_2, i.e. form tetrahedra with x_2 and the current interface triangles, and then use x_2 to triangulate its clause cupola if necessary (Figure 3.17.d.). This is possible by part (b) of Condition 5. We continue by forming tetrahedra using z_T as apex (Figure 3.17.e, possible by Condition 5, part (d)). At last, we will include the triangle (c_{2C}, z_L, z_B). All these triangles are visible by part (d) of Condition 6. After all these tetrahedra are added the interface is ready for the next variable.

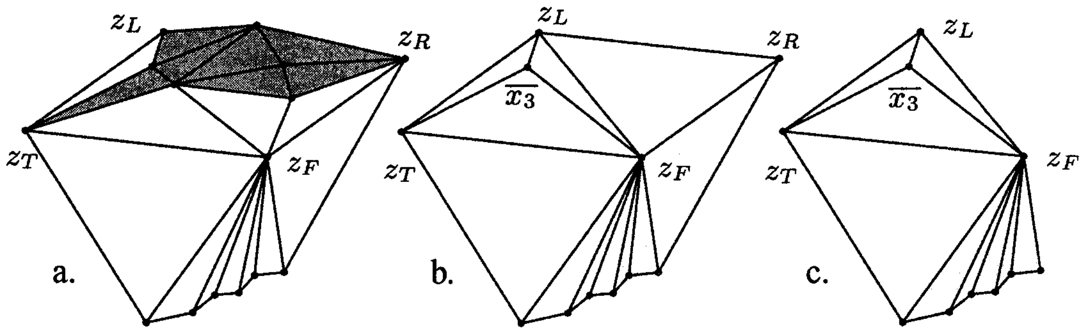

Figure 3.18: *The sweep for* $X_i = false$

Let us now consider the triangulation step in the case X_i is set to be $false$: We use the vertex z_F to triangulate the variable cupola as well as seven faces of the roof (see Figure 3.18): $(z_T, \overline{x_3}, z_A)$, $(\overline{x_3}, z_A, z_B)$, $(\overline{x_3}, z_L, z_B)$, (z_B, z_A, x_2), (z_B, x_2, z_R), (z_A, x_2, x_1), (x_2, x_1, z_R). The reader can see that on the roof we are leaving only the vertex $\overline{x_3}$. Next the tetrahedron (z_F, z_L, z_R, c_{2C}) is cut out. Hereby the interface triangle (z_F, z_R, c_{2C}) is replaced by (z_F, z_L, c_{2C}) (Figure 3.18.c.). Then $\overline{x_3}$ will be used as apex with the triangles in the interface. If the negated literal $\overline{X_i}$ is used to satisfy its clause j, the jth clause cupola is triangulated by $\overline{x_3}$. The interface advances as in the $true$-case. Then z_T can be used to form tetrahedra with the triangles in the interface. In the end the interface is again ready for the next variable.

How many tetrahedra can such a triangulation have? Triangulating all cupolas with a vertex in their visibility cones yields at most $(3m + 16)(C + V)$ tetrahedra (Proposition 3.6). In one step of the sweeping triangulation the tops of the roofs are each triangulated using six or seven tetrahedra (if the variable is unnegated or negated, resp.). Furthermore, the interface is triangulated by some vertices three times (in the positive case by x_1^i, by x_2^i, and by z_T^i) or two times (in the negative case by $\overline{x_3^i}$ and by z_T^i). The interface contains in each step between C and $2C$ triangles. Eventually, in either case there is one more tetrahedron (see above). An upper bound for the size of this triangulation is therefore

$$\begin{aligned}\#T &\leq (3m+16)(C+V)+7V+3CV+1 \\ &= m(3C+3V)+\underbrace{16C+23V+3CV+1}_{p_T(V,C)}\end{aligned}$$

What is the number of the vertices of the logical polytope in terms of the number of clauses and variables? We have V logical variables and C clauses in the SAT instance. We have m interior points each of the vertex-edge chains we added (later we will determine the value of m as a polynomial function of V and C). We observe that we have $3m+6$ vertices in each cupola, hence we have $(3m+6)(V+C)$ for all cupolas. We have in each roof nine vertices, two of them are shared with the subsequent roof except for the last roof. Hence the total number of vertices in roofs is $7V+2$. We have left only the $2C+1$ vertices along the spine. In conclusion, the number of vertices of P is

$$\begin{aligned}n &= (3m+6)(V+C)+7V+2+2C+1 \\ &= m(3C+3V)+\underbrace{8C+13V}_{p_n(C,V)}+3\end{aligned}$$

We had said before that a "bad" triangulation (where at least one cupola skylight is triangulated by a vertex not lying in its visibility cone) has at least $n+m-3 = m(3C+3V+1)+p_n(C,V)$ tetrahedra. On the other hand a "good" triangulation has at most $m(3C+3V)+p_T(C,V)$ tetrahedra. We can now set $m = m(C,V) = p_T(C,V) - p_n(C,V) + 1$. Then, if a good triangulation exists, its size is smaller than or equal to $K = n+m-4$, and if not, all triangulations are larger than K. □

3.4 Constructions

So far we have only talked about the *properties* of the various construction steps, the properties of the Schönhardt polytope, the properties of the cupola and the properties of the logical polytope. In order to show the NP-hardness of MINTRIANG(3) however, we have to demonstrate how to *construct* each of these entities. We split this into several steps: In Section 3.4.2 we will show how to construct a Schönhardt polytope and an m-cupola if only the rest of the polytope and a visibility cone is known. Then in Section 3.4.3 we will show how to construct visibility cones satisfying certain conditions. This in turn we will use in the last step in Section 3.4.4 when we construct the logical polytope itself.

Many times we will need to construct a point or a hyperplane satisfying certain "open" polynomial conditions, i.e. its coordinates must satisfy certain strict polynomial inequalities. Our construction paradigm is it to start in a special position (if we want to construct a point beyond a face, we start on the face) and perturb this point slightly. This perturbation has to be carried out in a way that the encoding length of the coordinates of the points do not get to large. Section 3.4.1 will describe the algorithmic version of this paradigm.

3.4.1 Constructing Points Satisfying Open Conditions

For our constructions we will often have situations where we want to move points from a special position to a more general position while other conditions are still satisfied. As long as these other conditions are *open* conditions, i.e. polynomial strict inequalities in the coordinates of the points, we can use the following paradigm:

Elementary steps of construction include operations such as taking the join of two or three points, intersecting planes and lines, putting points on polynomial curves, etc. The coordinates of the resulting construction elements are therefore polynomials in coordinates of the input elements. On the other hand, we will have requirements on the positions of the points with respect to some planes or other points on lines etc. All these conditions can be formulated as strict polynomial inequalities in coordinates of the construction elements. An essential element of our construction is that our systems of strict polynomial inequalities will depend on *one single parameter* ϵ. All these polynomial inequalities are satisfied at $\epsilon = 0$, but an additional requirement for us is $\epsilon > 0$. The following lemma describes a polynomial algorithm to find a number ϵ_0 such that all $0 < \epsilon \leq \epsilon_0$ solve the inequality system.

Lemma 3.10 *1. Suppose* $p(\epsilon) = a_d\epsilon^d + \cdots + a_1\epsilon + a_0$ *is a polynomial with* $p(0) > 0$*. Let* $\epsilon_0(p) := \min\left(1, \dfrac{a_0}{2(|a_1| + \cdots + |a_d|)}\right)$*. Then for* $0 \leq \epsilon \leq \epsilon_0(p)$ *we have* $p(\epsilon) > 0$*.*

Hence, the construction of ϵ_0 *can be done in time polynomial in the encoding length of the coefficients of p, and* ϵ_0 *has polynomial encoding length.*

2. $p_1, \ldots, p_l$ *are univariate polynomials such that* $p_1(0) > 0, \ldots, p_l(0) > 0$ *then there is a rational number,* $\epsilon_0 > 0$*, such that* $p_1(\epsilon) > 0, \ldots, p_l(\epsilon) >$

0 *for all* $0 < \epsilon \leq \epsilon_0$. *Moreover, the encoding length of* ϵ_0 *is bounded by a polynomial in the encoding length of the coefficients of* $p_1, \ldots, p_l$.

PROOF: For $0 \leq \epsilon \leq 1$ we have that $a_i \epsilon^i \geq -|a_i|\epsilon$. The reason is that for $a_i \geq 0$, $a_i\epsilon^i \geq 0 \geq -|a_i|\epsilon$, and for $a_i < 0$, $a_i\epsilon^i > a_i\epsilon = -|a_i|\epsilon$. Hence for $0 \leq \epsilon \leq \epsilon_0(p)$

$$p(\epsilon) \geq \sum_{i=1}^{d} -|a_i|\epsilon + a_0 > -\sum_{i=1}^{d} |a_i| \frac{a_0}{2\sum_{i=1}^{d}|a_i|} + a_0 > 0.$$

For the second part, take the value $\epsilon_0(p_1, \ldots, p_r) := \min(\epsilon_0(p_1), \ldots, \epsilon_0(p_r))$. Now *all* the conditions are simultaneously satisfied. □

For example in order to construct a point beyond some triangular facet we might want to take the barycenter of this face and move it just a bit out of the polytope such that all other facet-defining inequalities are still satisfied. The point then has the coordinates $p(\epsilon) = p_{bary} + \epsilon \cdot d_{normal}$. For $\epsilon = 0$ this point is in the polytope therefore satisfies all facet-defining inequalities $a_i^T p(\epsilon) - b_i > 0$. The lemma says that we can find an ϵ_0 such that for all $0 < \epsilon \leq \epsilon_0$ the point $p(\epsilon)$ is beyond the facet. In our proof of the hardness of MINTRIANG we will also require polynomiality of the encoding length of ϵ_0. It is guaranteed.

Of course, in general the real solutions of a multivariate system of inequalities coming imposed by geometric requirements may be empty, but our steps of construction reduce everything to sequentially solving easy univariate systems of inequalities.

3.4.2 Constructing a Cupola from a Visibility Cone

In this section we will show that cupolas can be attached to any face of a frame polytope using intermediate polytopes and that the visibility cone can be prescribed. The following theorem does not have the full strength we need for the construction of the "logical polytope" we need to show the hardness of MINTRIANG. Later, we will use a slightly stronger version which we will present at the end of this section. However, this theorem captures the main ideas used to construct a cupola.

Lemma 3.11 (Cupola Construction from a Given Visibility Cone) *Let* F *be a facet of a* 3*-polytope* P*, and* V *be a triangular cone such that* $F \cap V$ *is a* *triangle in the relative interior of F, and m be a positive integer. Then there is an m-cupola C beyond F of P such that P is beyond (A_1, A_2, A_3) of C and such that V is the visibility cone of C. Moreover, the input length of C is polynomial in the input lengths of P, V and m.*

Before we come to the proof, we will exhibit a necessary condition of the visibility cone V of a cupola C and the facet the cupola is being glued upon. It will imply that we cannot directly attach a cupola to a face (as in [51]), but we have to construct an intermediate polytope first.

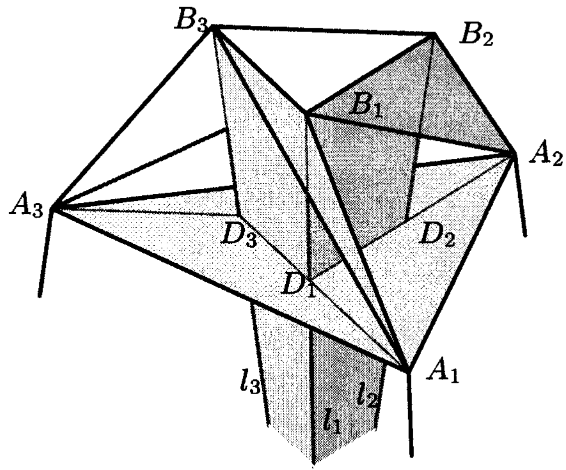

Figure 3.19: *Collinearity condition in the base triangle of a cupola*

Lemma 3.12 *Let A_1, A_2, A_3, B_1, B_2, B_3 be vertices in Schönhardt position. Define l_1 to be the intersection line of planes $B_3B_1A_1$ and $B_1B_2A_2$, lines l_2 and l_3 are defined accordingly (Figure 3.19, note that they contain the extreme rays of V). The lines l_1, l_2, and l_3 intersect the relative interior of the bottom face (A_1, A_2, A_3) of a cupola C. The intersection points D_1, D_2, and D_3 are forced to have the following collinearities: $A_1D_1D_2$, $A_2D_2D_3$, and $A_3D_3D_1$.*

PROOF.l_1 enters the Schönhardt polytope S in point B_1, runs along facet (A_1, B_1, B_3) until it reaches the edge (A_1, B_3) where it goes into the interior of S. Then the relative interior of (A_1, A_2, A_3) contains the point D_1. In this way, D_1, D_2, A_2 are all on the planes $A_1A_2A_3$ and $B_1B_2A_2$. □

Proof of Lemma 3.11: We proceed in three steps. The lines l_1, l_2, l_3 are defined as in Lemma 3.19.

The bottom triangle (A_1, A_2, A_3). We will now construct an intermediate polytope beyond F which will have a triangular facet (A_1, A_2, A_3) which is (1) parallel to F, and which is (2) intersected by the cone V in a triangle (D_1, D_2, D_3) in the relative interior such that (3) the collinearity condition from Lemma 3.12 holds.

To do this, we place a plane H parallel to and slightly above F such that the intersection points D_i of H and l_i $(i = 1, 2, 3)$. Also H has to be so close to F that the l_i do not cross between H and F. By prolonging the line segment D_3D_1 slightly beyond D_1 (staying in $P \setminus F$) we obtain point A_1, analogously construct A_2 and A_3 (Figure 3.20). Taking the convex hull of F and the points A_1, A_2 and A_3 gives then the intermediate polytope, whose face (A_1, A_2, A_3) has the collinearity condition. These constructions are polynomially constructible in the sense of Lemma 3.10.

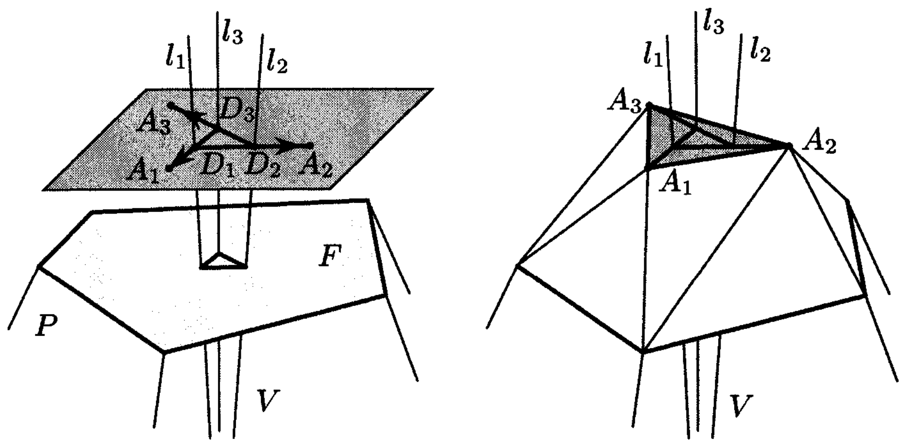

Figure 3.20: *Building the intermediate polytope for the cupola*

The frame of the cupola. As in the construction of the bottom facet (A_1, A_2, A_3), we place a plane H' parallel and slightly above this facet. The intersection of H' and V is the triangle (B_1, B_2, B_3) (B_1 is on the same extreme ray of V as D_1 and so on). See Figure 3.21.

It is clear from the construction that triangles (D_1, D_2, D_3), (A_1, A_2, A_3), and (B_1, B_2, B_3) are parallel and all oriented the same way. Therefore it is not hard to check that the points A_1, A_2, A_3, B_1, B_2, and B_3 are vertices of a Schönhardt polytope whose visibility cone is V. Polynomiality of this part of the construction follows from Lemma 3.10 as well.

Attaching the vertex-edge chains. Now that the frame of a cupola is done, i.e., the vertices $A_1, \ldots, B_3$ are in Schönhardt position, it remains to patch

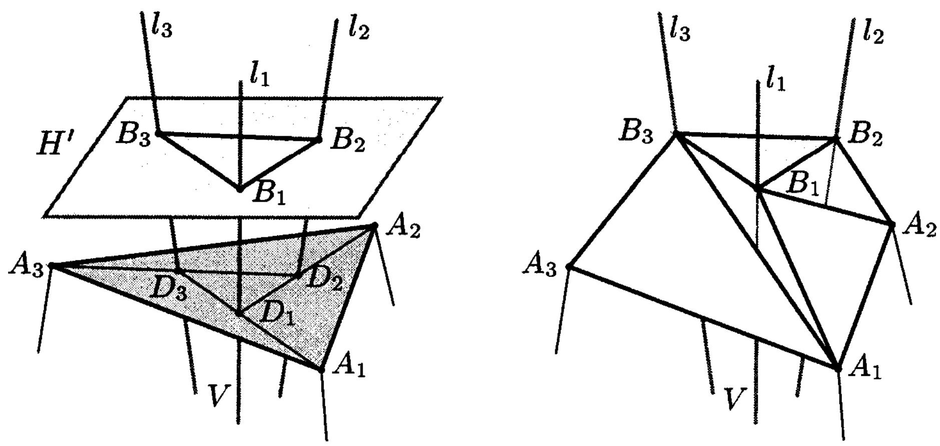

Figure 3.21: *Building the frame of a cupola*

the key structures of Lemma 2.2, the vertex-edge chains q_i^j ($i = 1, \dots, m$, $j = 1, \dots, 3$), to the sides of the frame $conv(P \cup \{A_1, \dots B_3\})$.

Given triangular faces (a, q_0, q_{m+1}) and (b, q_0, q_{m+1}) of a convex polytope P and a plane G which (strictly) separates points q_0 and q_{m+1}. We claim that we can construct points $q_1, \dots, q_m$ beyond the edge (q_0, q_{m+1}) of P such that the convex hull of $P \cup \{q_i\}$ has the properties of Lemma 2.2 and such that the points $q_1, \dots, q_m$ lie on the same side of G as q_0. Moreover, the input length of the constructed points is polynomially bounded in the input length of P and G.

By applying our claim three times, we will conclude our proof. The vertices q_i^j are placed beyond edge (A_j, B_{j+1}), vertices B_j and A_{j+1} take the roles of a and b, G is the plane spanned by B_j, A_{j+1} and B_{j+2}. It is easy to check that this is exactly what we want for Lemma 2.2 and for the cupola conditions.

Now we prove the claim. We will put the points q_i ($i = 1, \dots, m$) on a parabola segment, beyond the edge (q_0, q_{m+1}). Let H be a plane containing q_0 and q_{m+1} which also intersects the interior of P. This plane has the property that it contains points beyond edge (q_0, q_{m+1}). It is constructible in polynomial time. Let v be the sum of the two normal vectors of planes aq_0q_{m+1} and bq_0q_{m+1}, and H the plane containing q_0 and q_{m+1} parallel to v.

Let now D be the intersection point of G and (q_0, q_{m+1}). Let w be a vector of direction of the intersection line of G and H, such that starting at D it is pointing out of P. Now let $E_\varepsilon = D + \varepsilon w$ for $\varepsilon > 0$ to be specified later. For small ε, E_ε is beyond (q_0, q_{m+1}). Hence the parabola defined according to Lemma 3.13, stated and proved below, by $p(0) = q_0$, $p(1/2) = E_\epsilon$, and

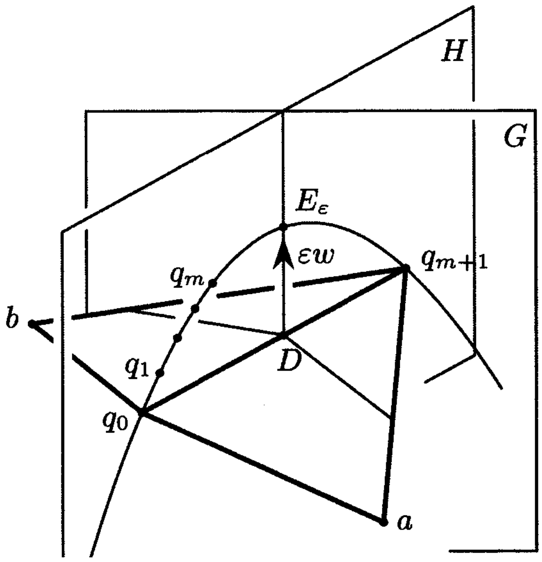

Figure 3.22: *Construction of the vertex-edge chain*

$p(1) = q_{m+1}$ lies entirely in H, and for arguments between 0 and 1 passes just beyond (q_0, q_{m+1}). Let $q_i = p(i/(4m))$ for $i = 1, \ldots, m$. For small ε all those points are beyond (q_0, q_{m+1}) and on the same side of G as q_0 (polynomial conditions, use Lemma 3.10). Also, they are in convex position such that the convex hull of $P \cup \{q_1, \ldots, q_m\}$ has exactly the required face lattice. □

Lemma 3.13 *Let p_0, p_1, p_2 be three non-collinear points in $\mathbb{R}^3$ and t_0, t_1, t_2 be three distinct real numbers. Then there is a unique curve $p : \mathbb{R} \to \mathbb{R}^3$ such that $p_0 = p(t_0)$, $p_1 = p(t_1)$, and $p_2 = p(t_2)$ which is quadratic in every coordinate. Furthermore, all points on $p(t)$ are in the plane spanned by p_0, p_1, and p_2, and they are in convex position. Also a plane containing $p(r)$ and $p(l)$ for some $r \neq l$ which does not contain all of p has all points between l and r on one of its sides and all other points on the other side.*

PROOF.Since p_0, p_1, p_2 have to be on the t_0, t_1, t_2 positions of the curve

$$p(t) = \begin{pmatrix} a_x + b_x t + c_x t^2 \\ a_y + b_y t + c_y t^2 \\ a_z + b_z t + c_z t^2 \end{pmatrix},$$

we have the condition

$$\begin{pmatrix} a_x & b_x & c_x \\ a_y & b_y & c_y \\ a_z & b_z & c_z \end{pmatrix} \begin{pmatrix} 1 & 1 & 1 \\ t_0 & t_1 & t_2 \\ t_0^2 & t_1^2 & t_2^2 \end{pmatrix} = \begin{pmatrix} \vdots & \vdots & \vdots \\ p_0 & p_1 & p_2 \\ \vdots & \vdots & \vdots \end{pmatrix}.$$

By the non-singularity of the Vandermonde matrices, there is a unique solution to $a_.$, $b_.$, $c_.$ given the p_i and t_i.

The curves which are quadratic in every coordinate are linear transforms of the moment curve $m(t) = (1, t, t^2)$. This curve lies entirely in the $x = 1$ plane, is convex, and has the condition that it intersected by each plane at most twice (or it is in this plane). All these properties are invariant under linear transformations. □

3.4.3 Constructing a Visibility Cone

In order to use the cupola as a basic building block, we need to have a visibility cone that contains a specified set of vertices and intersects the relative interior of some face. Once we have that we can construct the cupola as described in the previous section. The set will consist of all vertices lying in a specified plane.

Lemma 3.14 *Let H be a plane which intersects the relative interior of some face F of a polytope P, and let $S = \{v_1, \ldots, v_s\}$ be the set of vertices of P lying in H, not including the vertices of F. Let $S' = \{w_1, \ldots, w_{s'}\}$ be a set of points in relint$(F) \cap H$. It is possible to construct a triangular cone V which intersects F in a triangle that lies in the relative interior of F and V contains S and S' in its interior and no other vertex of P.*

The reader may not see at this point the purpose of the set S', but we will justify it at the end of this section.

PROOF.$P \cap H$ is a polygon. Without loss of generality, $F \cap H$ is horizontal and situated on the top of the polygon $P \cap H$ (see Figure 3.23). Let l be the line connecting the leftmost point of S' and leftmost vertex of S (the one encountered first when walking around $P \cap H$ counterclockwise, starting at $F \cap H$). Analogously, let r be the line connecting M and the rightmost vertex of S.

The area between l and r (in H) is already a cone containing S and no other vertices of P. We will perturb it in a way that the other conditions are satisfied as well.

First shift l and r parallely outwards, guaranteeing that they still intersect $F \cap H$ in its relative interior (easy open conditions); we obtain l' and r'. Also, let f' be a line in H parallel to F just outside P, i.e., such that l' and r' intersect f' in the same order as $F \cap H$ (again using Lemma 3.10).

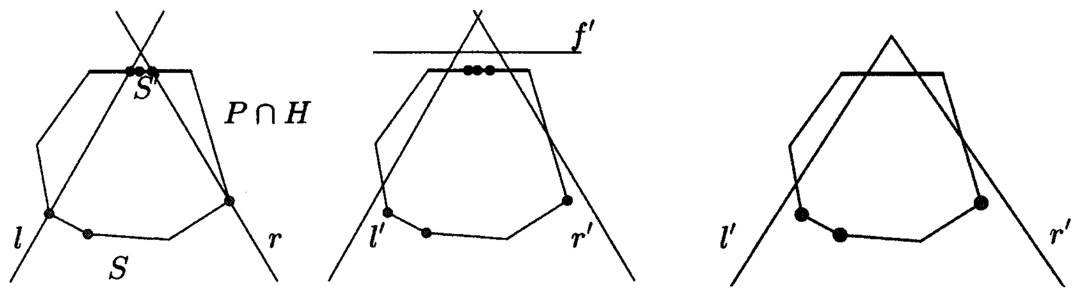

Figure 3.23: *Construction of l and r, then l' and r' (viewed in H)*

Now we will rotate H about l' and r' and f', getting three planes bounding the desired triangular cone: Let H be oriented in some way, and $a_H x \geq b_H$ be its defining inequality. Let v be some point which lies on the positive side of H. Let $G_{l'}$ be the plane through l' and v. By construction, all vertices in S lie on the same side of $G_{l'}$, so we can orient it such that S is on its positive side. Let $a_{l'}x \geq b_{l'}$ be its defining inequality. Perform the same construction for r' and f' obtaining $G_{r'}$ and $G_{f'}$, also orienting them in a way that v is on their respective positive sides. Let $G^{\varepsilon}_{l'}$ be the plane defined by $(a_H + \varepsilon a_{l'})x \geq b_H + \varepsilon b_{l'}$. This plane contains l' and for small ε it is very close to H. Hence, it is the *rotation of H about l' in the direction of plane $G_{l'}$*. Also let $G^{\varepsilon}_{r'}$ be defined by $(a_H + \varepsilon a_{r'})x \geq b_H + \varepsilon b_{r'}$, and $G^{\varepsilon}_{f'}$ be defined by $(-a_H + \varepsilon a_{f'})x \geq -b_H + \varepsilon b_{f'}$.

Obviously, all points in S and in S' are on the positive sides of the planes $G^{\varepsilon}_{l'}$, $G^{\varepsilon}_{r'}$, and $G^{\varepsilon}_{f'}$. For small $\varepsilon > 0$, these planes do not "sweep" over vertices of P which are not in S, and it is easy to see that in this case, there are no vertices of P that satisfy all three new inequalities. Also for small ε, the points in F satisfying all three inequalities define a triangle in the relative interior of F with endpoints $G^{\varepsilon}_{l'} \cap G^{\varepsilon}_{r'} \cap F$, $G^{\varepsilon}_{r'} \cap G^{\varepsilon}_{f'} \cap F$, and $G^{\varepsilon}_{f'} \cap G^{\varepsilon}_{l'} \cap F$. Hence, the set of all points satisfying the three inequalities is a triangular cone V with the desired properties. The conditions on ε are open polynomial conditions according to Lemma 3.10. □

This lemma can be used to build *one* cupola over the facet F. However, there might be problems if we keep on constructing around the polytope, like adding more cupolas over other facets of P. The visibility cone we just constructed

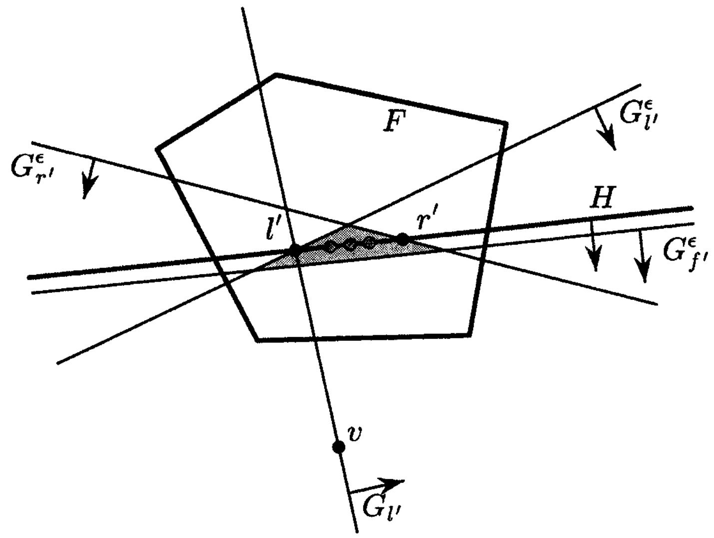

Figure 3.24: *Rotated hyperplanes, viewed by their intersections with F*

might "catch" points we construct later. But these constructions all happen *beyond* facets of P, so we can use the following lemma to construct all cupolas one after the other without their visibility cones catching extra vertices.

Lemma 3.15 *Let $H_1, \dots, H_n$ hyperplanes, intersecting facets $F_1, \dots, F_n$ of a polytope P with the restriction that $F_i \cap H_j = \emptyset$ for all $i \neq j$. Then $P_{\text{beyond } F_i} \cap H_j = \emptyset$ for all $i \neq j$.*

PROOF.Assume there is a point u in $P_{\text{beyond } F_i} \cap H_j$ ($i \neq j$). Then this point also lies in $(P \setminus F_i) \cap H_j$, but on the non-positive side of F_i. Let v be a point in $F_j \cap H_j$, Then v is also in $(P \setminus F_i) \cap H_j$ (since $F_j \subseteq P \subseteq P \setminus F_i$), but on the positive side of F_i. Hence, there must be a point w on the line segment $[u, v]$ which is on the hyperplane containing F_i. The whole segment lies in $P \setminus F_i$, hence every point on it has to satisfy all of P's defining inequalities except that of F_i. So w lies *in* the facet F_i. But it also lies in H_j (the whole line segment does), which contradicts the assumption $F_i \cap H_j = \emptyset$. □

In Section 3.1 we will need an additional condition: Given a set of lines in the plane H (of Lemma 3.14) that pierce the face F, we want to be sure that these lines also pierce the skylight of the constructed cupola. (This condition will play an important role when we want to force so-called blocking conditions, see Section 3: At some point two tetrahedra spanned by two skylights and two respective visible vertices v and v' will have to intersect in their interiors. This is already guaranteed if the corresponding lines g and g' intersect inside the polytope.)

The next theorem specifies the way in which we will use all the preceding lemmas in our construction in Section 3:

Theorem 3.16 (Full-Strength Cupola Construction) *Let H_i be planes that intersect facets F_i of a polytope P in their relative interiors such that $F_i \cap H_j = \emptyset$ for all $i \neq j$. Let $S_i = \{v_1^i, \ldots, v_{s_i}^i\} := (vert(P) \cap H_i) \setminus F_i$, and $L_i = \{g_1^i, \ldots, g_{s_i'}^i\}$ sets of lines. Assume further that each of the lines g_j^i pierces the relative interior of F_i and is incident to some v_k^i.*
Then we can sequentially construct all cupolas C_i beyond the faces F_i such that in the resulting polytope their visibility cones contain S_i and no other vertices. In addition, the skylight of the cupola over each F_i is pierced by the lines in L_i.

PROOF.The theorem follows from the ideas in Lemmas 3.11 and 3.14. In the construction of the visibility cone over facet F_i, we invoke Lemma 3.14 with the polytope $P \cup \bigcup_{j \neq i} P_{\text{beyond } F_j}$. The set S_i' is of course $\{l \cap F_i | l \in L_i\}$. The cupola construction was such that the cupolas over F_j were always *beyond* the facet F_j, so the constructed visibility cones contain no vertices of the other cupolas. In order to have the lines in L_i pierce the skylight of cupola i we have to alter the construction of the cupola in Lemma 3.11: when we put the planes parallel to F, we do it in such a way that the triangles (A_1, A_2, A_3) and then (B_1, B_2, B_3) are pierced by these lines. These are both open conditions on the distance of the planes to F. $\square$

3.4.4 Constructing a Logical Polytope

Proof of Lemma 3.7: The construction will be carried out in five stages. By the time we end the construction all five requirements of the definition of logical polytopes must be satisfied, but three of the conditions will not be met until the last stage.

1. Give coordinates of the basic wedge, with rectangular faces on top for each variable.

2. Attach the roofs for each variable, giving preliminary coordinates for the literal vertices and preliminary coordinates for the points on the lower edge (the spine of the wedge).

3. Perturb the literal vertices to their final positions.

4. Perturb the vertices on the spine of the wedge.

5. Attaching the variable cupolas following the procedures of Section 2.

In every step we will build a construction element (a point, a line, or a plane) whose coordinates are polynomials in the construction elements up to that particular moment. Hence, the encoding length of each new construction element is bounded by a linear function of the encoding length of the construction so far. The number of construction steps is polynomially bounded in C and V. Hence the encoding length of the whole construction is also polynomially bounded in C and V. Note however, that the coordinates themselves will in general be exponentially large.

Instead of writing explicit (and highly cumbersome) coordinates for the construction elements, we rely on Lemma 3.10 to ensure that such coordinates can be found if one has really the desire to see a particular logical polytope. A key property of Stages 2–4 in the construction is that the geometric conditions we want to determine a finite collection of strict polynomial inequalities *in a single variable.* Then, by Lemma 3.10, we know there is an appropiate polynomial size solution. In subsequent stages of the construction similar new systems, for other independent parameters, will be solved, preserving what we had so far, but building up new properties.

Stage 1: The basic wedge. Consider the triangular prism which is the convex hull of the six points $c_0 = (0,0,0)$, $c_{2C} = (0,1,0)$, $z_T^V = (0,0,1)$, $z_F^1 = (1,0,1)$, $z_L^V = (0,1,1)$, and $z_R^1 = (1,1,1)$. See Figure 3.25(a). In order to obtain a convex structure on the top of the wedge, we consider the function $f(x) = x(1-x)+1$. The vertices of each roof boundary (that is z_T^i and z_F^i as well as z_R^i and z_L^i) will lie on the surface $z = f(x)$. More specifically, $z_F^i = z_T^{i+1} = (i/V, 0, f(i/V))$ and $z_R^i = y_L^{i+1} = (i/V, 1, f(i/V))$ for $i = 0, \ldots, n$. By the concavity of f, the points are indeed in convex position and their convex hull, the *wedge* has the desired face lattice (see Figure 3.25(b)).

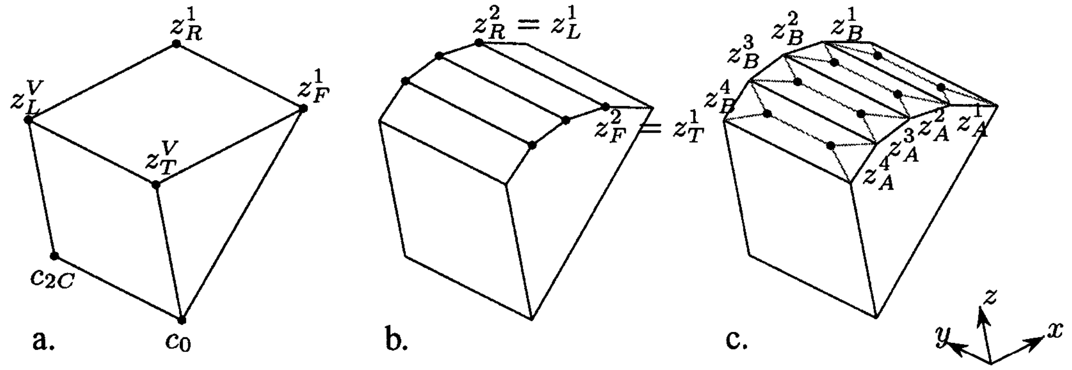

Figure 3.25: *Construction of the wedge*

So far none of the conditions we want are satisfied (not even partially).

Stage 2: The roofs. We will first attach the points z_A^i and z_B^i to the quadrilateral face $(z_L^i, z_R^i, z_T^i, z_F^i)$. Then we give preliminary coordinates to the literal vertices and to the vertices on the spine.

Let $z_A^i = 1/2 \cdot (z_T^i + z_F^i) + (0, 1/3, t_{\text{roof}})$ and $z_B^i = 1/2 \cdot (z_T^i + z_F^i) + (0, 2/3, t_{\text{roof}})$ where t_{roof} is a non-negative parameter that is called the *roof height*. That is the points have the same x coordinate as the midpoint between z_T^i and z_F^i, y coordinate $1/3$ and $2/3$ respectively, and height t_{roof} over the face $(z_T^i, z_F^i, z_L^i, z_R^i)$. We want to choose t_{roof} in a way that z_A^i and z_B^i are beyond the facet $(z_T^i, z_F^i, z_L^i, z_R^i)$ (see Figure 3.25(c)). We can easily achieve this by the technique presented in Lemma 3.10: The only possibly concave edges are the (z_T^i, z_L^i). One restriction is therefore that all determinants $\det(z_T^i, z_L^i, z_A^{i-1}, z_A^i)$ have to be positive. These are finitely many open quadratic conditions on t_{roof}. For $t_{\text{roof}} = 0$ the points z_A^i an z_B^i are inside the facets $(z_T^i, z_F^i, z_L^i, z_R^i)$, hence the edges in question are trivially convex. We will get more polynomial constraints on t_{roof} below and then solve all simultaneously to find the suitable roof height.

The spine of the wedge is still a line. We now put preliminary points $c_0, \ldots, c_{2C}$ on this line. Let

$$u(j) = \frac{1}{2}\frac{j}{2C}$$

and $c_j = (0, u(j), 0)$ for $j = 0, \ldots, 2C-1$, and $c_{2C} = (0, 1, 0)$ (see Figure 3.26). As an auxiliary point, let b_l be the barycenter of the points c_{2l-2}, c_{2l-1}, and c_{2l} ($l = 1, \ldots, C$). At this moment, this point $b_l = c_{2l-1}$. Later, as we perturb the spine vertices b_l will move accordingly, always $b_l = 1/3(c_{2l-2} + c_{2l-1} + c_{2l})$.

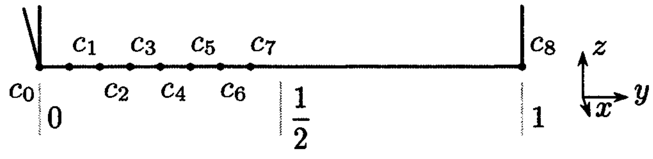

Figure 3.26: *Preliminary coordinates for the spine vertices*

Now we want to give initial positions to the literal vertices. Say variable X_i occurs unnegated in clauses l_1 and l_2 and negated in l_3. Note that l_j depend on the variable we are considering. For instance, in our example logical formula on p. 45, for variable X_1, $l_1 = 1$, $l_2 = 3$, and $l_3 = 2$. But for variable X_2, $l_1 = 2$, $l_2 = 3$, and $l_3 = 1$.

The preliminary literal vertex x_1^i is the intersection of the $y = u(2l_1 - 1)$ plane with the line connecting z_F^i and z_B^i. We do the same for the other positive occurrence of X_i and obtain the preliminary x_2^i. For the negative occurrence of X_i, we take the line connecting z_T^i and z_B^i, intersect it with the $y = u(2l_3 - 1)$ plane, and obtain the preliminary $\overline{x_3^i}$. We join the preliminary x_1^i and b_{l_1} by a line d_1^i (this line lies in the $y = u(2l_1 - 1)$ plane). Do the analogue process for x_2^i and $\overline{x_3^i}$, obtaining d_2^i and d_3^i. Later we will move the vertices $x_1^i, x_2^i, \overline{x_3^i}$ along their respective lines d_1^i, d_2^i, d_3^i a little out of polytope in order to turn them into extreme points. The lines d_j^i will also be used for blocking conditions.

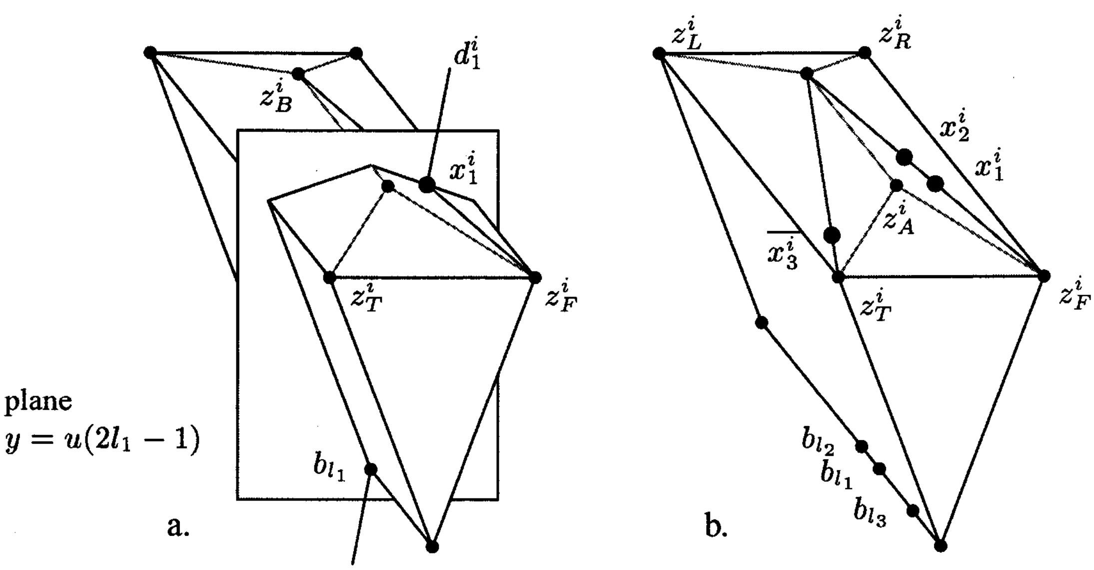

Figure 3.27: *Construction of the literal vertices in the X_i slice of the wedge*

Let H^i be the plane that contains z_T^i and z_F^i and the midpoint of the edge (z_L^i, z_B^i) (Figure 3.28). The only vertices above H^i are $x_1^i, x_2^i, \overline{x_3^i}, z_A^i$, and z_B^i, and the only vertices *on* H^i are z_T^i and z_F^i. This follows from the convexity of the current polytope.

Let g_1^i (g_2^i) be the line in the plane H^i which is incident to z_F^i and intersects the line d_1^i (d_2^i). Note that this intersection point lies in the segment (x_1^i, c_{2l_1-1}) (the line segment (x_2^i, c_{2l_2-1})), thus in the interior of the constructed polytope. Analogously, let g_3^i be the line in the plane H^i which is incident to z_T^i and intersects the line segment $(\overline{x_3^i}, c_{2l_3-1})$. It can be verified that if the roof height is small (z_L^i, z_R^i, z_B^i) is pierced by the g_j^i in its relative interior. This is another strict polynomial inequality in t_{roof}. It will be the planes H^i and lines g_j^i ($i = 1, \ldots, V$) from which we make the visibility cones for the cupolas of

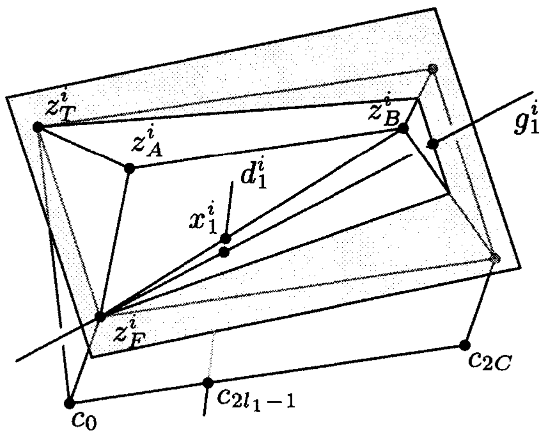

Figure 3.28: *Construction of H^i and g_1^i*

variables X_i according to Theorem 3.16.

It is important to note right now that the non-blocking conditions are satisfied for this special position of the vertices. We do not want the tetrahedron $(z_T^i, z_L^i, z_R^i, z_B^i)$ and the triangle $(x_1^i, c_{2l_1-2}, c_{2l_1})$ to intersect . From this we get strict polynomial inequalities on t_{roof}. They are satisfied for $t_{\text{roof}} = 0$ since the y coordinates of the spine vertices c_l are smaller than $1/2$. A suitable value of t_{roof} can be found solving the univariate inequality system we accumulated in our discussion (Lemma 3.10). It is easy to check that the sweeping conditions are also satisfied for the preliminary position of the points $x_1^i, x_2^i, \overline{x_3^i}$. So far we have met two of the five required conditions to have a logical polytope.

Stage 3: Literal vertices Now we put the final x_j^i $(j = 1, 2, 3)$ a little outward on line d_j^i (Figure 3.27). A little for x_1^i and x_2^i means that the positive literal vertices lie in a plane parallel to the face $(z_R^i, z_B^i, z_A^i, z_F^i)$ very close to it. We treat $\overline{x_3^i}$ similarly. If the three literal vertices are moved a sufficiently small distance t_{literal} the face lattice of what we get after taking the convex hull is as Figure 3.13 (a) in all roofs. See also the Schlegel diagram in 3.13 (b).

By construction H_i contains z_F^i and z_T^i, and the $y = u(2j-1)$ planes contain all literal vertices corresponding to clause j. This will become important for the visibility conditions (see Stage 5). Also, for small t_{literal} the non-blocking and sweeping conditions are satisfied.

Although we do not have the blocking condition yet auxiliary lines can be set up: As above, let l_1, l_2, l_3 be the clauses to which the literal vertex x_1^i, x_2^i, $\overline{x_3^i}$ belong. We made sure that the line segments (c_{2l_1-1}, x_1^i) and (z_F^i, z_B^i) intersect in their respective relative interiors. Hence, by the construction of line g_1^i, it is also pierced by (x_1^i, c_{2l_1-1}) between z_F and the face (z_L^i, z_R^i, z_B^i). (Analogously, (c_{2l_2-1}, x_2^i) and $(z_F^i, g_2^i \cap (z_L^i, z_R^i, z_B^i))$ as well as (c_{2l_3-1}, x_3^i) and $(z_T^i, g_3^i \cap (z_L^i, z_R^i, z_B^i))$ intersect in their relative interiors). Later on this intersection will evolve into the real blocking conditions using Theorem 3.16.

Stage 4: The perturbing the vertices on the spine of the wedge. We now perturb the points c_j on the spine of the wedge. Every even-indexed c_{2l} is changed to lie on a parabola, and for the moment the odd-indexed vertices c_{2l-1} are changed to lie on the line connecting c_{2l-2} and c_{2l}. The y coordinates of all points stay the same:

$$c_{2l} = \left(\frac{1}{2}(y-1)^2 \cdot t_{\text{even}}, y, (y-1)^2 \cdot t_{\text{even}}\right).$$

Note that by the $1/2$ in the x coordinate, the points are moved *into* the polytope. The changes (parameter t_{even}) must be small enough that the convex hull now has the desired appearance (Figure 3.29) and the non-blocking conditions and the sweeping conditions are still satisfied. Once more we appeal to Lemma 3.10. The polynomials inequalities are now on the variable t_{even} and the sweeping and non-blocking were satisfied at $t_{\text{even}} = 0$. The reader should note that while the constructed vertices in the roofs do not change coordinates, dependent construction elements like the lines d_j^i (connecting x_j^i and c_{2l_j-1}) and g_j^i (lying in H^i and intersecting d_j^i) change when the spine vertices move. However, the parameter t_{even} has to be small enough that the preliminary blocking conditions are still met: g_j^i still pierce the facet (z_L^i, z_R^i, z_B^i) in its relative interior, and g_j^i and d_j^i intersect in the interior of the polytope.

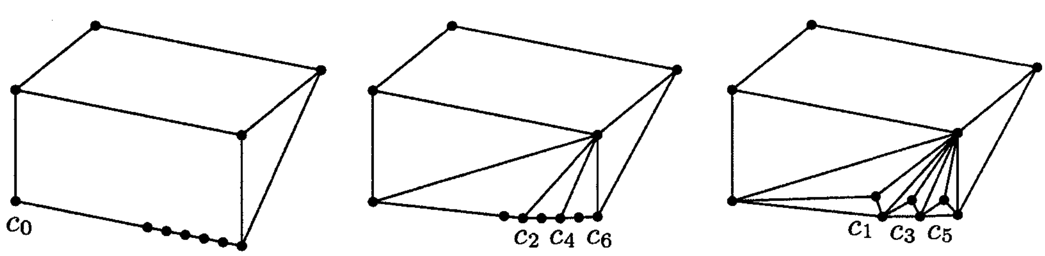

Figure 3.29: *Perturbation of the vertices on the spine*

Now we move the odd points c_{2l-1} beyond the face $G_l = (c_{2l-2}, c_{2l}, z_T^0)$: to this end, we choose a point p_l beyond G_l and move to $c_{2l-1} + t_{\text{odd}}(p_l - c_{2l-1})$. Such a point p_l is easily found by taking a normal to G_l through its barycenter and moving outwards while staying beyond the face (note that this involves again Lemma 3.10, see the definition of *beyond*). The parameter t_{odd} is chosen small enough: Convexity and the correctness of the face lattice are easily achieved. Also the sweeping conditions are valid for slight moves. Keeping t_{odd} small also guarantees the non-blocking conditions: the tetrahedron $(x_1^i, c_{2l_1-2}, c_{2l_1-1}, c_{2l_1})$ is only slightly bigger than just the triangle $(x_1^i, c_{2l_1-2}, c_{2l_1})$ which did not intersect the tetrahedron $(z_T^i, z_L^i, z_R^i, z_B^i)$ (x_2^i and $\overline{x_3^i}$).

For the blocking conditions, let X_i be the jth logical variable in clause l. Note that now the line d_j^i intersects the triangle $(c_{2l-2}, c_{2l-1}, c_{2l})$ in its relative interior. The lines g_j^i are updated as the lines d_j^i move. Since t_{odd} is small, g_j^i still pierces the facet (z_L^i, z_R^i, z_B^i) in its relative interior, and g_j^i and d_j^i intersect in the interior of the polytope. Note that d_j^i is still in the $y = u(2l-1)$ plane because the y coordinates of the spine vertices were conserved.

Stage 5: Attaching the cupolas. It remains to construct all the cupolas. Over the facets (z_L^i, z_R^i, z_B^i) $(i = 1, \ldots, V)$ we construct cupolas using the planes H_i and sets of lines $\{g_1^i, g_2^i, g_3^i\}$, and over the facets $(c_{2l-2}, c_{2l-1}, c_{2l})$ $(i = 1, \ldots, C)$ we construct the clause cupolas using the $y = g(2l-1)$ planes and the sets of lines $\{d_j^i \mid X_i\text{'s } j\text{th occurence is in clause } l\}$. We invoke Theorem 3.16 and get the final polytope. By this construction, it is convex, has the correct face lattice, and the visibility conditions are satisfied.

The reader will recall that g_j^i and d_j^i intersect in the interior of the polytope. Say again variable X_i occurs unnegated in clauses l_1 and l_2 and negated in l_3. By Theorem 3.16 g_j^i pierces the skylight of the cupola corresponding to variable X_i and d_j^i pierces the skylight corresponding to its clause l_j. Hence, the tetrahedron spanned by z_F^i and the variable X_i's skylight together with the tetrahedron spanned by x_1^i (x_2^i) and clause l_1's skylight (l_2's skylight) intersect in their interiors. Analogously, the tetrahedron spanned by z_T^i and the variable X_i's skylight and the tetrahedron spanned by $\overline{x_3^i}$ and clause l_3's skylight intersect in their interiors. These are exactly the blocking conditions.

All other conditions concerned only points we constructed before, so they are still satisfied. The final polytope is therefore a logical polytope. □

3.5 Minimal Triangulation with Interior Points

We consider now again the concept of triangulations using additional interior points. We want to construct a polytope now whose minimal triangulation uses at least k interior points. Then we will show that MINTRIANGIP$(3, l)$ is NP-hard.

Theorem 3.17 *Given two numbers $h \geq 1$, $k \geq 1$ there is a simplicial convex 3–polytope P such that every triangulation of P using less than h interior points has at least k tetrahedra more than a triangulation of P with h suitably chosen interior points.*

PROOF: Let P be a simplicial 3-polytope on n_P vertices with at least h facets. For h of these (triangular) facets $F = (A_1, A_2, A_3)$ we construct a visibility cone as follows:

Choose three points D_1, D_2, D_3 in F such that A_i, D_i, and D_{i+1} are collinear, but such that D_1, D_2, and D_3 are not collinear. Then consider a point C beyond F and a plane H such that the intersections $B_i = H \cap D_iC$ are still beyond F, but on the other side of C w.r.t. D_i. These constructions are easy to accomplish (even in polynomial time by Lemma 3.10). See Figure 3.30.

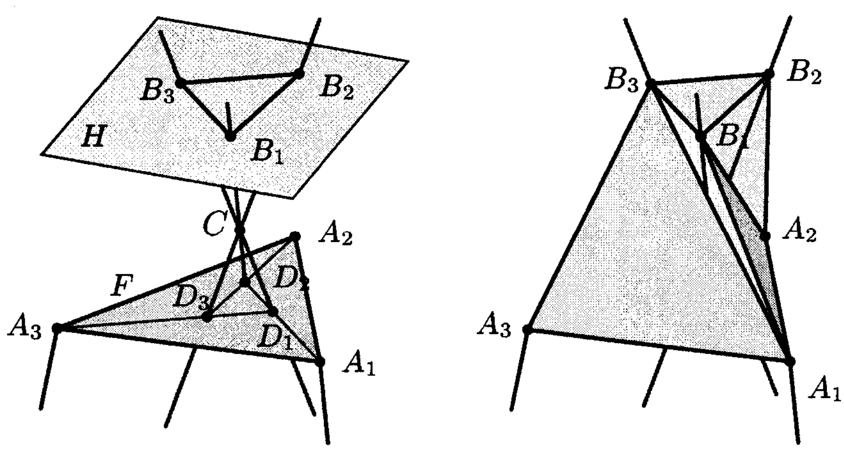

Figure 3.30: *Constructing disjoint visibility cones over many facets*

It is not hard to see that the points A_i, B_i $(i = 1, \ldots, 3)$ are in Schönhardt position and that the visibility cone is the triangular cone with apex C and rays CB_i. Note that this visibility cone contains no other points of P. Now we can complete the cupola construction (according to the construction in Lemma 3.11) and get an m-cupola where m is to be specified later.

We perform this construction for all h triangular facets. Then the subsets of the visibility cones consisting of the points *below* the corresponding skylight do not intersect each other and contain no vertices.

Let T be a triangulation of this polytope P_h^{IP} with less than h interior points. We will show that this is a large triangulation: There is one visibility cone which contains no interior point and no vertex of P_h^{IP}. Hence by Proposition 3.5 T has at size at least $n + m - 3$ where n is the number of vertices of P_h^{IP}. By construction $n = n_P + 3h(m + 1)$.

We will now construct a small triangulation using h interior points: Of course we will place one interior point in each visibility cone. Let C' be a point in the visibility cone of a cupola which is very close to C. With very close we mean close enough that it can triangulate all facets of the Schönhardt polytope from *within*. This point is easy to find since C can see the lower facets (A_i, A_{i+1}, B_i) and the facet (A_1, A_2, A_3) from within, and any point in the visibility cone can see all other facets from within. Then we can triangulate the Schönhardt polytope using 8 tetrahedra. The vertex-edge chains can be triangulated using $3(m + 1)$ tetrahedra. The original polytope P remains, by a pulling triangulation we can triangulate it using at most $2n_P - 7$ tetrahedra. We obtain a triangulation of P_h^{IP} with at most $2n_P - 7 + 8h + 3h(m + 1)$. Hence for $m \geq k+n+8h-4$ this triangulation is at least k tetrahedra smaller than any triangulation using less than h interior points. $\square$

Corollary 3.18 MINTRIANGIP$(3, h)$ *is NP-hard for all* $h \geq 0$. *Thereby h is* not *part of the input.*

PROOF: Let P_f be the logical polytope constructed from a given SAT formula f *without* the vertex-edge chains. Beyond the front triangle of the wedge we construct a vertex-edge chain consisting of h vertices. (We have to take care that they are not in any visibility cones, but this is an open condition satisfied by c_0, so if we build everything very close to c_0 we are fine. By Lemma 3.10 this can be achieved in polynomial time.) Over the triangles of

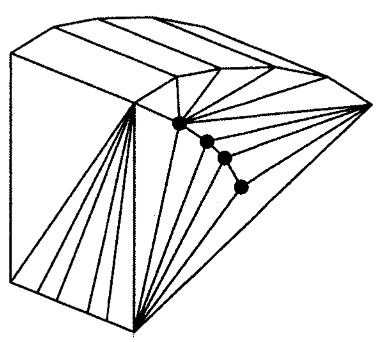

Figure 3.31: *Constructing h triangular facets where we can glue cupolas with empty visibility cones*

these triangles we can construct Schönhardt polytopes with empty and disjoint visibility cones as in the previous proof. Now construct all vertex-edge chains (in the logical polytope and in the h Schönhardt polytopes) with length $m = 5h+V+3+p_T(C,V)-p_n(C,V)$. The resulting polytope has $n = 3m(C+V+h)+4h-1+p_n(C,V)$ vertices. We have to show now that f is satisfiable if and only this polytope has a small triangulation (less than $n+m-3$ tetrahedra) using at most h interior points.

In a small triangulation of less than $n+m-3$ tetrahedra all skylights must be triangulated by points in their visibility cone. In particular the h new skylights must be triangulated. But their visibility cones contain no vertices, so we have to introduce interior points as above. Furthermore, they do not intersect, so we must place one interior point per new cupola. This leaves no interior points, so the remaining cupolas must be triangulated with vertices of the polytopes. By the properties of the logical polytope this means that the formula f is satisfiable.

On the other hand if f is satisfiable we triangulate all new cupolas with interior points as in the previous lemma, needing $8+3(m+1)$ tetrahedra per cupola. Then we can triangulate the the new structures we attached to the front face of the wedge by forming the $h-1$ tetrahedra of two consecutive new vertices and the edge (c_0, z_F^1). It remains the wedge plus a pyramid over its front face. We triangulate the front face by pulling c_0 and using this triangulation (having V triangles) we pull the last new vertex. We are left with the logical polytope which we have seen we can triangulate using $m(3C+3V)+p_T(C,V)$. Summing up all numbers we get

$$\#T = 3m(C+V+h)+9h-1+V+p_T(C,V).$$

By construction of m this number is smaller than $n+m-3$. □

Part II

Combinatorial Polytopes and Finding Maximal Triangulations

Chapter 4

Realizations of Polytopes

The face lattice of a polytope (the partially ordered set of all faces) contains the *combinatorial* information of the polytope. It is interesting to ask how much of the information about the polytope is already stored in the face lattice.

Let us first look at an example where the face lattice does not contain enough information to give a valid answer for every realization, namely maximal triangulations. The 3-dimensional cube can be realized as the unit cube spanned by the vertices $\{0,1\}^3$; its maximal triangulation has 6 tetrahedra[1]. However, another realization of the cube is an upright prism over a non-rectangular four-gon whose maximal triangulation has 7 tetrahedra (Figure 4.1).

We want to answer the following questions:

1. Is there a polytope such that in every realizations a face has a certain shape? In other words, can we *prescribe* the exact shape of a face using only the combinatorial data of the polytope?

2. Is it hard to find maximal boundary triangulations over *all realizations* of a polytope?

The answer to both questions is yes, and in a very strong sense:

Yes, for every d-polytope $\boldsymbol{G}$ with algebraic coordinates we can construct a $d{+}2$-polytope $\boldsymbol{P}$ such that any realization of $\boldsymbol{P}$ has a face which is projectively equivalent to $\boldsymbol{G}$. This result is best possible in two ways: We cannot hope to

[1] This is easy to see: Each tetrahedron has volume at least $1/3!$, so in order to fill the body of volume 1 there can be at most 6; this bound is achieved by any coning triangulation.

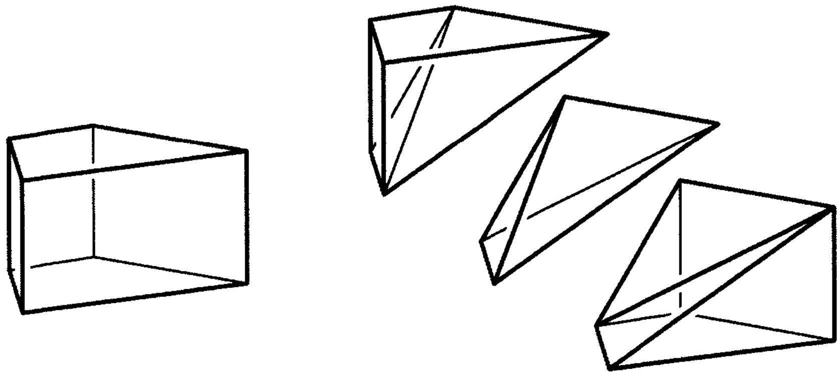

Figure 4.1: *The maximal triangulation of the prism over a non-rectangular square has 7 tetrahedra: As opposite edges of this polytope are not parallel like in the unit cube, they span a tetrahedron; the rest of the polytope can be triangulated using three tetrahedra on each side of this tetrahedron.*

accomplish more than projective equivalence, because projective images of a polytope have the same face lattice. And we cannot hope to prescribe the exact shape of polygons with non-algebraic coordinates (more precisely, polygons for which all projective images have non-algebraic coordinates); this is an implication of the Tarski-Seidenberg theorem [40, 53].

And yes, it is as hard as the existential theory of the reals to find the maximal boundary triangulation over all realizations of d-polytopes for $d \geq 5$. This implies for instance NP-hardness.

One of the main (outside) tools for these results is Richter-Gebert's Universality Theorem [46]: He showed that realization spaces of 4-polytopes can be as complicated as any primary semialgebraic sets[2] can be. Viewed this way the combinatorial structure already encode pretty much. He encoded polynomials equation and inequality systems, in short *polynomial systems* into 4-polytopes which then had the following properties:

1. in each realization there is a face which "encodes" a solution to the polynomial systen and

2. for each solution is a realization with a face which "encodes" this realization.

[2] a basic primary semialgebraic set is the solution set of a system of polynomial equations and *strict* inequalities with algebraic coefficients

Since we rely heavily on his construction we recall it in great detail. Along with introducing the precise notation and stating general-knowledge prerequisites this is the main contents of Chapter 4.

In Chapter 5 we will use Richter-Gebert's construction to come up with polytopes which prescribe the exact shape of one of its faces. First we will show how to prescribe 2-faces of 4-polytopes by making explicit use of Richter-Gebert's construction. Then we will prescribe d-faces of $d + 2$-polytopes for $d \geq 3$. For this we will employ the results of the first part of this section.

Chapter 6 is dedicated to the proof of the hardness of finding maximal boundary triangulations over all realizations. We will encode a given polynomial system (a so-called *driving system*) in a polytope: the polytope will have a realization with a large boundary triangulation if and only if the driving system has a solution.

Most facets of this polytope have only triangulations which exists in every realization and are easy to compute. (For an intuition, think pyramids over polygons: all of their triangulations are present in every realization.) But the shape of one facet is prescribed to have only two possible realizations, one with a large triangulation and another with a smaller triangulation. The driving system is now encoded in such a way that the realization with the large triangulation is only possible if the driving system has a solution.

Overview of this Chapter

In this chapter we want to introduce all the techniques and all notation we need for our constructions and which are either common knowledge or other peoples results. Among the first are the remarks in Section 4.1 about combinatorial vs. realized polytopes. In Section 4.2 we define the concept of a combinatorial polytope which *prescribes* a property of one of its faces. We give various constructions which are examples for this concept and which are prerequisites of the Universality Theorem for polytopes. Actually, all these constructions can be found in Richter-Gebert's monograph [46] in one way or the other. In Section 4.3 we state the Universality Theorem in a form which we can use in later chapters. We give some, not all, details of the construction, just so much that the reader can appreciate our changes to it in the said later chapters. At the end of this chapter, in Section 4.4, we have placed an appendix on projective spaces and projective transformations and how polytopes fit into this picture. The results there are again common knowledge, nevertheless we added the proofs to keep this treatment self-contained.

4.1 Combinatorial Polytopes and Realizations

Faces of Polytopes

A polytope is equivalently the convex hull of a finite number of (finite) points and the intersection of a finite number of halfspaces such that the result is bounded.

The faces of a polytope are the intersection of the polytope with a supporting hyperplane (a hyperplane which defines a halfspace which together with the hyperplane itselft contains the whole polytope). Equivalently it is the zero set of an *affine functional* (a linear functional followed by a scalar translation) which is non-negative on the whole polytope. By this definition the empty set and the whole polytope are also (so-called improper) faces. Notice that faces are again polytopes. The dimension of a face is the dimension of the containing affine space. If the dimension of a polytope is d then we call the 0-dimensional, 1-dimensional, $d-2$-dimensional, and $d-1$-dimensional faces vertices, edges, ridges, and facets, respectively. We will call the set of all vertices of a polytope $\boldsymbol{vert}(\boldsymbol{P})$, and the set of all facets $\boldsymbol{facets}(\boldsymbol{P})$.

Labeling of Faces

If two vertices p and q are connected by an edge we denote the edge by $p \vee q$. Similarly, if two edges e and f share a vertex we denote it $e \wedge f$. This is an abuse of the notation for the join ($\vee$) and the meet ($\wedge$) of projective subspaces (see Section 4.4 for the definitions) since we will *use the name of a face and the name of its projective closure interchangably*. This extends to the case where $\text{conv}(\{a, b, c, d\})$ is a quadrangular face of a polytope $\boldsymbol{P}$; then we will also denote it by the name of the plane $a \vee b \vee c$ — and actually mean the intersection of the plane with the polytope $(a \vee b \vee c) \cap \boldsymbol{P}$. Also we will use the name of a flat *piercing* a polytope and the set of intersection of the flat and the polytope interchangably: If p and q are vertices not connected by an edge, then $p \vee q$ denotes both the line containing both of them as well as its segment inside of the polytope, i.e. the diagonal connecting p and q.

Let us see this notation at work in polygons. If the edges of a polygon are labeled by an index set $X = (a_1, \ldots, a_n)$ we will refer to this polygon by $G(X) = G(a_1, \ldots, a_n)$ (see Figure 4.2). Very often we will use the integers $1, \ldots, n$ as *edge* labels. The vertices of the polygon are then the intersection $i \wedge i+1$ of consecutive edges (addition modulo n the number of edges). The diagonals or *chords* of the polygon we label particularly: (i, j) is the

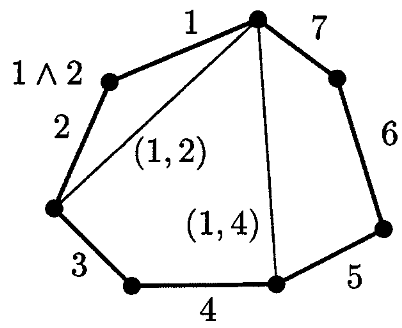

Figure 4.2: *The polygon* $G(1,2,3,4,5,6,7)$

diagonal spanned by the vertex before edge i and the vertex after j, i.e.$(i,j) = (i-1 \wedge i) \vee (j \wedge j+1)$. Naturally, we identify (i,i) and i.

Schlegel Diagrams

We remind the reader of an important tool to visualize higher-dimensional polytopes, the *Schlegel diagrams*. Schlegel diagrams are constructed using central projections of polytopes onto one of their facets. The projection central is chosen beyond that facet so that all other polytope vertices are projected *into* the chosen facet (see Figure 4.3 and [61]). The images of the $d-2$-dimensional faces (ridges) on the chosen facet subdivide this facet cells: these are the images of the other facets and they are actually projectively equivalent to them.

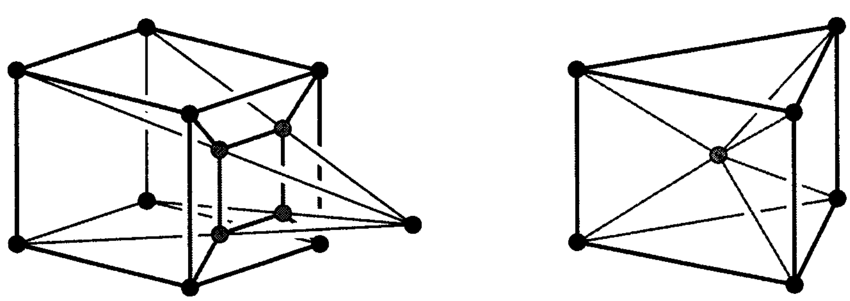

Figure 4.3: *Left: Construction of Schlegel diagram of a* 3*-dimensional cube; Right: Schlegel diagram of a pyramid over a* 3*-dimensional prism (*4*-polytope)*

Combinatorial Polytopes

The *face lattice* of a polytope is the partially ordered set of all faces of a polytope, ordered by inclusion. The face lattice captures the combinatorial structure of its boundary. It is an atomic lattice, i.e. each face is completely determined by its vertex set (which is the set of atoms beneath it in the ordering). But it is also coatomic, i.e. each face is the intersection of the facet that include it. Hence the face lattice is already completely determined by the vertex sets of its facets. This gives rise to the following definition.

Definition 4.1 *The* combinatorial polytope of a polytope $\boldsymbol{P}$ *is the set* P *of vertex label lists for all facet of* $\boldsymbol{P}$*, i.e. the set* $P = \{vert(\boldsymbol{F}) | \boldsymbol{F} \text{ facet of } \boldsymbol{P}\}$.

$\boldsymbol{P} =$

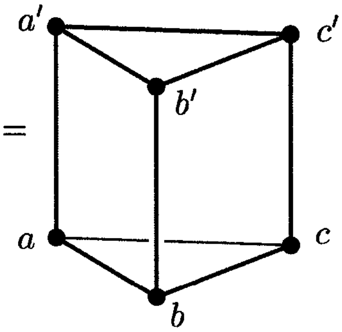

$$P = \{ \; \{a,b,c\}, \{a',b',c'\}, \{a,b,a',b'\}, \{b,c,b',c'\}, \{a,c,a',c'\} \; \}.$$

Figure 4.4: *Example of a polytope and its combinatorial polytope*

Consider two polytopes $\boldsymbol{P}$ and $\boldsymbol{Q}$ which share a facet $\boldsymbol{F}$ such that $\boldsymbol{P}$ and $\boldsymbol{Q}$ are on different sides of $\boldsymbol{F}$ and such that $\boldsymbol{P} \cup \boldsymbol{Q}$ is a convex polytope. The facets of the resulting polytope are of course all facets of $\boldsymbol{P}$ and all facets of $\boldsymbol{Q}$ except the facet $\boldsymbol{F}$, so the combinatorial polytope of $\boldsymbol{P} \cup \boldsymbol{Q}$ is easily described: all facets of $\boldsymbol{P}$ and all facets of $\boldsymbol{Q}$, except the common facet. We use this operation as a generalization of the concept of combinatorial polytope in the following recursive definition:

Definition 4.2 *A set* P *of lists of vertex labels is called a* (general) combinatorial polytope *on a vertex set* X

1. *if it is the combinatorial polytope of a polytope* $\boldsymbol{P}$ *with* $\mathbf{vert}(P) = X$*, or*

2. *if there are combinatorial polytopes* Q *and* Q' *on vertex label sets* Y *and* Y' *such that* $X = Y \cup Y'$*,* $Y \cap Y' \in Q \cap Q'$*, and* $P = Q \cup Q' \setminus Y \cap Y'$.

In the latter case, if we denote $F = Y \cap Y'$, we say P is the connected sum *of Q and Q' along facet F. We write $P = Q\#_F Q'$.*

(See Section 4.2.2 for more details on connected sums.)

The *face lattice* of a combinatorial polytope P is the set of all intersections of the sets in P, partially ordered by inclusion. This lattice is still atomic and coatomic. Hence the elements of the face lattice can again be thought of as combinatorial polytopes, these are the (combinatorial) *faces* of P. The maximal vertex sets will of course correspond to the (combinatorial) *facets* of the combinatorial polytope.

Once a combinatorial polytope P is fixed, by an abuse of notation we will often identify the vertex sets of a face (the elements of the face lattice P) with the face itself (which is again a combinatorial polytope). For instance the combinatorial polytope $F = \{\{a, b\}, \{b, c\}, \{a, c\}\}$ is a facet of the polytope P in Figure 4.4 and is identified with its vertex set, the element $\{a, b, c\}$ of the face lattice of P. Also the intersection of F with the face $G = \{\{a, b\}, \{a, a'\}, \{b, b'\}, \{a', b'\}\}$ we will write $F \cap G = \{a, b\} \equiv \{\{a\}, \{b\}\}$.

A polytope $\boldsymbol{P}$ is a called the realization of a combinatorial polytope P if their face lattices are isomorphic, i.e. if there is a vertex labeling of $\boldsymbol{P}$ such that the vertex sets on each facet of $\boldsymbol{P}$ are exactly the facets of P. A combinatorial polytope is *realizable* if it has a realization. Notice that all faces of combinatorial polytopes are realizable. Also in our constructions in Chapters 5 and 6 all combinatorial polytopes will be realizable.

Notice that if a combinatorial polytope is realizable its (combinatorial) faces are in one-to-one correspondence with the (realized) faces of the realizing polytope, and of course they have the same vertex sets. The dimension of a (combinatorial) face F is then the length of a chain in the face lattice down to $\emptyset$. (This definition makes sense since face lattices of combinatorial polytopes is graded.)

Two (combinatorial or realized) polytopes P and P' are said to be *combinatorially equivalent* if there is a bijection f of their vertex sets such that F is a facet (face) of P if and only if $f(F)$ is a facet (face) of P'. Two projectively equivalent polytopes are of course combinatorially equivalent, but the converse does not hold at all (it is already false for pentagons).

Examples: Pyramids and Prisms

Given a realizable combinatorial d-polytope P on a vertex label set X, the *pyramid* over P with apex $y \notin X$ is defined as the combinatorial $d + 1$-

polytope $\boldsymbol{pyr}(P,y) = \{F \cup \{y\} | F \in P\} \cup X$: Its facets are P itself and pyramids over the facets of P. We abbreviate $\boldsymbol{pyr}(P,y)$ by $\boldsymbol{pyr}(P)$ if we do not explicitly care about the label of the apex.

This way of giving a combinatorial definition for a special combinatorial polytope will occur often in the sequel. We must take care that the result really *is* a combinatorial polytope. In this case it is the *combinatorial polytope of a realized polytope*: Given any realization $\boldsymbol{P}$ of P in $\boldsymbol{R}^{d+1}$ and a point y outside of the d-dimensional hyperplane, the convex hull $\boldsymbol{conv}(\boldsymbol{P}, y)$ is a realization of $\boldsymbol{pyr}(P,y)$.

Let P and P' be two realizable d-polytopes on disjoint vertex index sets X and X' which are combinatorially equivalent. The *prism* over P and P' is defined as the combinatorial $d+1$-polytope $\boldsymbol{prism}(P,P') = \{X\} \cup \{X'\} \cup \{F \cup f(F) | F \in P\}$. (Obviously it is a combinatorial polytope.) The facets of this polytope are of its top and bottom faces which are combinatorially equivalent to P and prisms over all facets of P. In the case that $X' = \{x' | x \in X\}$ we abbreviate $\boldsymbol{prism}(P,P')$ by $\boldsymbol{prism}(P)$. See Figure 4.5.

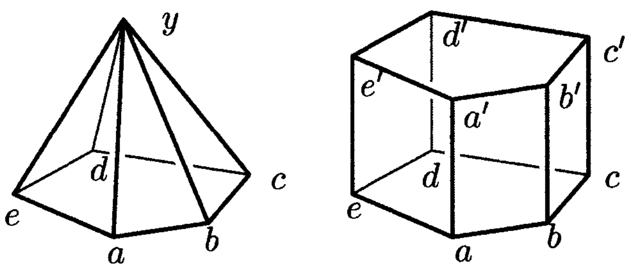

Figure 4.5: *Pyramid and prism over a pentagon P with vertex set* $\{a,b,c,d,e\}$

4.2 Prescribing Properties of Faces

In our constructions we want the combinatorial structure of our polytope to restrict the shape of some of its faces. The next definition encapsulates our need.

Definition 4.3 *Let F be a face of a combinatorial polytope P. Let E be a property that a realization of F can have or not have. We say that P* prescribes *E for F*

1. *if in every realization of P the realization of F has property E and*

2. *if every realization of F with property E can be completed to a realization of P.*

The main primitive construction of a polytope which prescribes the property of one of its faces is the so-called Lawrence extension which we will introduce presently.

Note that in the above definition we do not only demand that the face in every realization of the polytope has the desired shape, but also that if a realization of the face has the property that it can occur in a realization of the polytope. This will allow us to *superimpose* prescribed properties using connected sums (see Section 4.2.2).

This technique will be used to achieve the following goals: In Chapter 5, we will eventually be able to show that we can construct 4-polytopes ($d+2$-polytope) with a 2-face (d-face) which will be prescribed to be *projectively equivalent* to a given polygon (d-dimensional polytope). So the property E which we will prescribe is the projective equivalence of the face to the given polygon. In Chapter 6 we will prescribe weaker properties in order to show the hardness of finding maximal triangulations of polytope boundaries over all realizations.

4.2.1 Lawrence Extension

The Lawrence extension is a method to encode incidence information of a polytope in a polytope of one dimension higher. We will first give a description of its (realized) construction, then see some examples, and finally see what the properties are that Lawrence extensions prescribe.

Construction

Let $\boldsymbol{P}$ be a d-dimensional polytope and q a finite point *outside* of P. Embed this configuration into $\mathbb{R}^{d+1}$ such that $\boldsymbol{P}$ and q lie in a the affine closure $\overline{\boldsymbol{P}}$ of $\boldsymbol{P}$. Consider a line which intersects the hyperplane $\overline{\boldsymbol{P}}$ in the point q. Let q^+ and q^- be two distinct points on this line on the same side of $\overline{\boldsymbol{P}}$ such that q^- lies between q and q^+. The (realized) *Lawrence extension* is the polytope $\boldsymbol{\Lambda} = \boldsymbol{\Lambda}(\boldsymbol{P}, q) = \operatorname{conv}(\boldsymbol{P} \cup \{q^+, q^-\})$. Let $\Lambda(\boldsymbol{P}, q)$ be its combinatorial polytope.

Example: Tent

The easiest example of a Lawrence extension is the *tent*. This is a 3-dimensional (combinatorial) polytope consisting of a ground face and two additional vertices. Let e and f two non-adjacent edges of a polygon $\boldsymbol{P}$ and q the intersection $e \wedge f$ of the lines containing e and f. Then the tent is defined as $\boldsymbol{tent}^{e,f}(\boldsymbol{P}) := \Lambda(\boldsymbol{P}, q)$. See Figure 4.6.

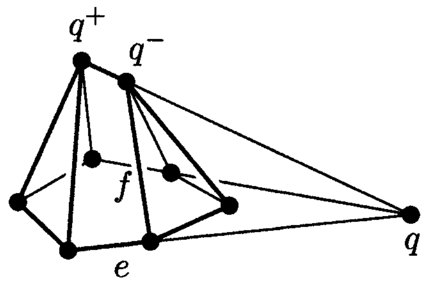

Figure 4.6: *The tent* $\boldsymbol{tent}^{e,f}(P)$

The Properties Prescribed by Lawrence Extensions

It is not hard to see that one facet of $\Lambda(\boldsymbol{P}, q)$ is $\boldsymbol{P}$ itself. If P is the combinatorial polytope associated to $\boldsymbol{P}$ then for every realization $\boldsymbol{P}'$ of P which occurs in a realization of $\Lambda(\boldsymbol{P}, q)$ there will be a point q' with the same relative position to the faces of $\boldsymbol{P}'$ as q to the faces of $\boldsymbol{P}$. In this way the Lawrence extension *prescribes* the existence of a point q with these properties on a combinatorial level.

Before we come to the precise statement, we need more notation. Define $\mathcal{F}^0(\boldsymbol{P}, q)$ the set of combinatorial facets of $\boldsymbol{P}$ whose realization in $\boldsymbol{P}$ has a supporting hyperplane containing q. Define also $\mathcal{R}^0(\boldsymbol{P}, q)$ the set of combinatorial ridges of $\boldsymbol{P}$ whose realization in $\boldsymbol{P}$ does not contain q in their affine hull, but has a supporting hyperplane containing q.

We say that a point q' *conforms* with a pair of sets $(\mathcal{F}^0, \mathcal{R}^0)$ with respect to a polytope $\boldsymbol{P}'$ if $\mathcal{F}^0 = \mathcal{F}^0(\boldsymbol{P}', q')$ and $\mathcal{R}^0 = \mathcal{R}^0(\boldsymbol{P}', q')$. Let $E^{\text{conform}}(\boldsymbol{P}, q)$ be the property of a realization $\boldsymbol{P}'$ of $\boldsymbol{P}$ that there is a point q' which conforms with $\mathcal{F}^0(\boldsymbol{P}, q)$ and $\mathcal{R}^0(\boldsymbol{P}, q)$.

Theorem 4.4 *The combinatorial polytope* $\Lambda(\boldsymbol{P}, q)$ *has a facet* P *which is the combinatorial polytope of* $\boldsymbol{P}$. *Furthermore,* $\Lambda(\boldsymbol{P}, q)$ *prescribes the property* $E^{conform}(\boldsymbol{P}, q)$ *for its face* P.

We will prove this result later in this Section on page 93.

Example: Transmitter polytope

Richter-Gebert uses Lawrence extensions in the proof of his Universality Theorem. One of the first examples is the *transmitter* polytope. It is the Lawrence extension $\Lambda(\boldsymbol{P}, q)$ where $\boldsymbol{P}$ is a prism over a polygon which is such that there is a point q which is the intersection of all lines containing edges which connect two corresponding points of the bottom and top face. Figure 4.7 shows the Schlegel diagram of the result.

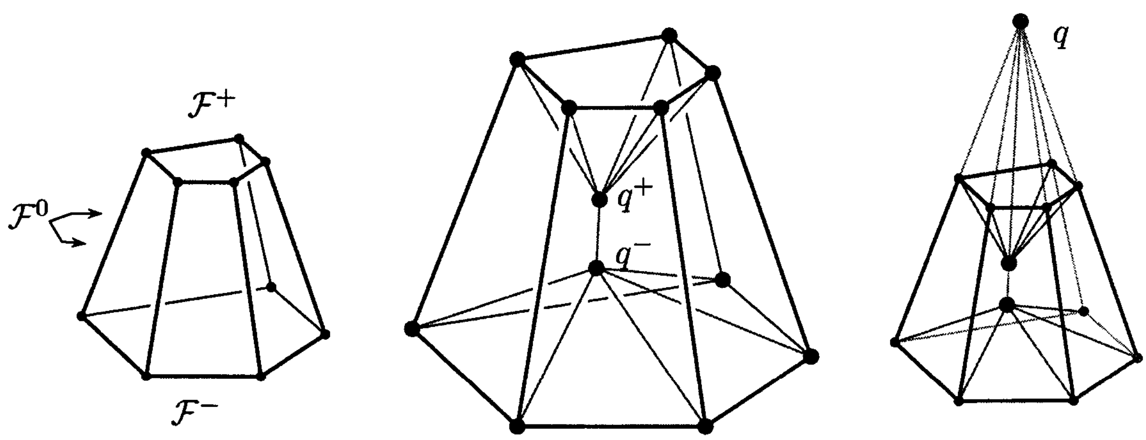

Figure 4.7: *The transmitter polytope*

Theorem 4.4 now implies that in every realization of this Lawrence extension there is a point q which lies in the planes defined by the quadrangular facets. (The set $\mathcal{F}^0$ contains all quadrangular facets on the sides of the prism.) Hence the lines supporting the edges of the prism connecting top and bottom polygons must all go through the point q: they are the intersections of two of these planes. But therefore the top and the bottom polygons are *projectively equivalent* since one is a central projection of the other, with q as projection center. It is for this projective equivalence that the polytope is named transmitter: it transmits information (the position of the points up to projective transformations) from one polygon to the other.

But prescribing a property had a second implication: any realization of the prism such that the edges connecting top and bottom polygons meet in a point can be completed to a realization of the transmitter polytope. But this means that the combinatorial structure of the transmitter does not depend on the shape of the ground face: any realization of the polygon with a parallel prism built over it can be completed to a realization of the transmitter, i.e. the transmitter prescribes nothing for the polygon (other than that it is a convex polygon).

This construction, i.e. taking the parallel prism over a polytope $\boldsymbol{P}$ and then Lawrence extension with the common direction of the parallel lines at infinity, canonically works also in higher dimensions.

Notice that even though the pyramids over the top and the bottom faces of the prism $\boldsymbol{P}$ are projectively equivalent in every realization of $\Lambda(\boldsymbol{P}, q)$, the top pyramid itself (and the bottom pyramid as well) is *not prescribed*, i.e. for any of its realization there is a realization of $\Lambda(\boldsymbol{P}, q)$ which contains this pyramid. This will become important when we use this polytope as "neutral" transmitter, i.e. a transmitter polytope which does not add any restrictions (see below).

Example: The polytope X, the Lawrence extension over a tent

We want to look at a 4-dimensional polytope that prescribes a property of a 2-face: Let $\boldsymbol{G} = \boldsymbol{G}(1,2,3,4,5,6)$ be a hexagon with the property that the lines containing the edges 2 and 5 as well as line containing the diagonal $(1,3)$ go through a point. Let $\boldsymbol{P} = \boldsymbol{tent}^{2,5}(\boldsymbol{G})$ be a tent over $\boldsymbol{G}$. Denote by a and b its vertices outside the ground face $\boldsymbol{G}$. The lines $(1,3)$ and $a \vee b$ both go through the point $2 \wedge 5$, so they span a 2-dimensional plane. The lines $(1 \wedge 6) \vee a$ and $(4 \wedge 5) \vee b$ also lie in this plane, so they meet in a point q. Let $\boldsymbol{X} = \boldsymbol{X}(1, \underline{2}, 3, 4, \underline{5}, 6)$ be the Lawrence extension $\boldsymbol{\Lambda}(\boldsymbol{P}, q)$. Notice that $\mathcal{F}^0$ consists of the triangular faces $a \vee 3$, $a \vee 4$, $b \vee 6$, and $b \vee 1\}$. $\mathcal{R}^0$ consists of the two edges 2 and 5.

The combinatorial polytopes X of $\boldsymbol{X}$ prescribes for the hexagon G that *the diagonal* $(1,3)$ *goes through the point* $2 \wedge 5$: If $\boldsymbol{X}$ is a realization of X then by the Lawrence extension (Theorem 4.4) the four triangular facets of the tent $1 \vee a$, $3 \vee b$, $4 \vee b$, and $6 \vee a$ meet in a point q. Hence the lines $(1 \wedge 6) \vee a$ and $(3 \wedge 4) \vee b$ also meet in q. Therefore the points a, b, $1 \wedge 6$, and $3 \wedge 4$ lie in a 2-dimensional plane H. Furthermore, the two quadrangular planes $2 \vee a \vee b$ and $5 \vee a \vee b$ and the ground face of the tent meet in a point, this must be the point $2 \wedge 5$. But since $2 \wedge 5$ lies on $a \vee b$, it must be in H. The intersection of H and the plane containing the ground face of the tent contains the points $1 \wedge 6$, $2 \wedge 5$, and $3 \wedge 4$. Hence these points are collinear. That we can extend a hexagon with the desired property to a realization of X follows directly from the construction.

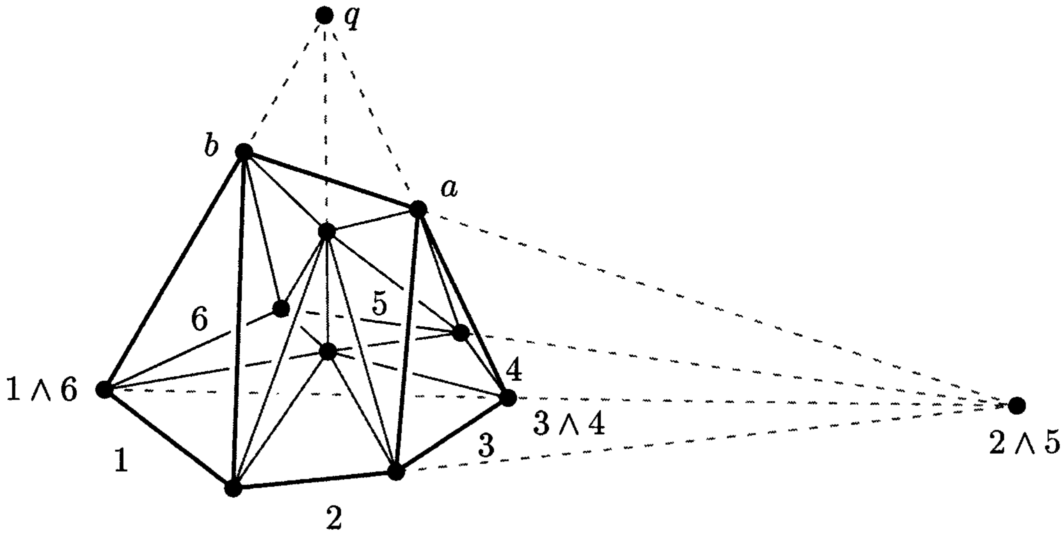

Figure 4.8: *The polytope* $X(1, \underline{2}, 3, 4, \underline{5}, 6)$

The Face Lattice of the Lawrence Extention

We will see now that the face lattice of a Lawrence extention is already determined of the position of the point q with respect to the facets of $\boldsymbol{P}$: first we will give the face latttice of $\boldsymbol{\Lambda}$ in terms of all faces, then (as a corollary) a description of its facets.

The point q defines three subsets of the set of faces of $\boldsymbol{P}$: the sets $\mathcal{G}^+ = \mathcal{G}^+(\boldsymbol{P}, q)$ and $\mathcal{G}^- = \mathcal{G}^-(\boldsymbol{P}, q)$ contain all faces of $\boldsymbol{P}$ with defining supporting hyperplanes having q on their *positive* side (*negative* side resp.). The set $\mathcal{G}^0$ is the set of all faces having a supporting hyperplane *containing* q. Note that all three sets contain the empty face and that $\mathcal{G}^0$ also contains the whole polytope $\boldsymbol{P}$ as a face.

Lemma 4.5 *The faces of* $\boldsymbol{\Lambda}$ *are:*

1. *all faces of* $\boldsymbol{P}$,
2. *conv*$(\boldsymbol{F} \cup \{q^+\})$ *for all faces* $\boldsymbol{F} \in \mathcal{G}^+$,
3. *conv*$(\boldsymbol{F} \cup \{q^-\})$ *for all faces* $\boldsymbol{F} \in \mathcal{G}^-$,
4. *conv*$(\boldsymbol{F} \cup \{q^+, q^-\})$ *for all faces* $\boldsymbol{F} \in \mathcal{G}^0$.

These are all faces. Hence the vertices of the polytope $\boldsymbol{\Lambda}$ *are the vertices of* $\boldsymbol{P}$ *and the points* q^+ *and* q^-.

PROOF: We will show first that all faces of $\boldsymbol{\Lambda}$ show up in the list we gave.

Let $\boldsymbol{F}$ be a face of $\boldsymbol{\Lambda}$ and H its supporting hyperplane (i.e. $\boldsymbol{\Lambda} \cap H = \boldsymbol{F}$). Call $\boldsymbol{F}_P$ the intersection of $\boldsymbol{F}$ with the hyperplane containing $\boldsymbol{P}$: $\boldsymbol{F}_P = \boldsymbol{F} \cap \overline{\boldsymbol{P}}$. We claim that $\boldsymbol{F}_P$ is a face of $\boldsymbol{P}$ and that

$$\boldsymbol{F} = \operatorname{conv}(\boldsymbol{F}_P \cup (\{q^+, q^-\} \cap \boldsymbol{F})),$$

i.e. $\boldsymbol{F}$ is the convex hull of $\boldsymbol{F}_P$ and the right choice of q^+ and q^-.

Notice that $\overline{\boldsymbol{P}}$ is a supporting hyperplane of $\boldsymbol{\Lambda}$ (q^+ and q^- are on the same side of $\overline{\boldsymbol{P}}$). Therefore $\overline{\boldsymbol{P}}$ defines the face $\boldsymbol{P}$ of $\boldsymbol{\Lambda}$. Hence $\boldsymbol{F}_P = (\boldsymbol{\Lambda} \cap \overline{\boldsymbol{P}}) \cap H = \boldsymbol{P} \cap H$ is a face of $\boldsymbol{P}$. The vertices of $\boldsymbol{F}$ come in two kinds: the ones that are in $\overline{\boldsymbol{P}}$, these are the vertices of $\boldsymbol{F}_P$, and the ones not in $\overline{\boldsymbol{P}}$, these are among $\{q^+, q^-\} \cap \boldsymbol{F}$. This shows the claim.

- Case 1. Neither q^+ nor q^- are in $\boldsymbol{F}$. Then $\boldsymbol{F} = \boldsymbol{F}_P$ is a face of $\boldsymbol{P}$.
- Case 2. The point q^+ is in $\boldsymbol{F}$, but q^- is not. Then $\boldsymbol{F}$ is the convex hull of $\boldsymbol{F}_P$ and q^+. Notice that since H contains q^- on its positive side and since q^- is between q and q^+, also q is on H's positive side. Hence $\boldsymbol{F}_P$ is a face in $\mathcal{G}^+$.
- Case 3. The point q^- is in $\boldsymbol{F}$, but q^- is not. Analoguously to case 2, $\boldsymbol{F} = \operatorname{conv}(\boldsymbol{F}_P \cup \{q^-\})$ and $\boldsymbol{F}_P \in \mathcal{G}^-$.
- Case 4. Both q^+ and q^- are in $\boldsymbol{F}$. Then $\boldsymbol{F}$ is the convex hull of $\boldsymbol{F}_P$ and q^+ and q^-. The hyperplane H contains q^+ and q^-, therefore also q. Hence $\boldsymbol{F}_P$ is in $\mathcal{G}^0$.

Conversely, it is easy to show that all listed sets are indeed faces. We will do this exemplarily for the set $\operatorname{conv}(\boldsymbol{F} \cup \{q^+\})$ for a face $\boldsymbol{F} \in \mathcal{G}^+$. Let H_F be a hyperplane that supports $\boldsymbol{F}$. The span $H = H_F \vee q^+$ is a d-dimensional hyperplane. It has q and all other vertices of $\boldsymbol{P}$ on one side of it. Since q^- is located between q and q^+, it also has q^- on the same side. Hence H is a supporting hyperplane of $\boldsymbol{\Lambda}$ and defines the face $\operatorname{conv}(\boldsymbol{F} \cup \{q^+\})$. □

This proof is an adaption of the proof of Lemma 3.3.3 of [46] which states the cone version of this lemma.

The facet description now follows easily: We let $\mathcal{F}^+ = \mathcal{F}^+(\boldsymbol{P}, q)$ and $\mathcal{F}^- = \mathcal{F}^-(\boldsymbol{P}, q)$ be the sets containing the facets which have q on their positive side (on their negative side, respectively). Then we weed out the faces of $\boldsymbol{\Lambda}$ which are not d-dimensional and obtain:

Corollary 4.6 *The polytope $\boldsymbol{\Lambda}$ has exactly the following facets:*

1. *$\boldsymbol{P}$ itself,*

2. *conv*$(\boldsymbol{F} \cup \{q^+\})$ *for all facets* $\boldsymbol{F} \in \mathcal{F}^+$,

3. *conv*$(\boldsymbol{F} \cup \{q^-\})$ *for all facets* $\boldsymbol{F} \in \mathcal{F}^-$,

4. *conv*$(\boldsymbol{F} \cup \{q^+, q^-\})$ *for all facets* $\boldsymbol{F} \in \mathcal{F}^0$, *and*

5. *conv*$(\boldsymbol{R} \cup \{q^+, q^-\})$ *for all ridges* $\boldsymbol{R} \in \mathcal{R}^0$.

The Proof of the Prescribability of q's Existence

PROOF OF THEOREM 4.4: In Corollary 4.6 we have shown that the Lawrence extension $\Lambda(\boldsymbol{P}, q)$ has a facet P which is the combinatorial polytope of $\boldsymbol{P}$. It remains to show

1. that for any realization $\boldsymbol{\Lambda}'$ of $\Lambda(\boldsymbol{P}, q)$ where $\boldsymbol{P}'$ is the realization of P there exists a point q' such that q' conforms with $\mathcal{F}^0$ and $\mathcal{R}^0$, i.e. $\mathcal{F}^0(\boldsymbol{P}', q') = \mathcal{F}^0(\boldsymbol{P}, q)$ and $\mathcal{R}^0(\boldsymbol{P}', q') = \mathcal{R}^0(\boldsymbol{P}, q)$ and

2. that all realizations $\boldsymbol{P}'$ with this property can be completed to a realization of $\Lambda(\boldsymbol{P}, q)$.

We begin by showing that in any realization $\boldsymbol{\Lambda}'$ of $\Lambda(\boldsymbol{P}, q)$ there is a point q' which conforms with $\mathcal{F}^0$ and $\mathcal{R}^0$. The vertices $q^{+'}$ and $q^{-'}$ of this realization must lie outside of the hyperplane containing the facet $\boldsymbol{P}'$. Let q' be the intersection point of the line $q^{+'} \vee q^{-'}$ and this hyperplane. Any supporting hyperplane H of a face $\boldsymbol{conv}(\boldsymbol{G}' \cup \{q^{+'}, q^{-'}\})$ for $G \in \mathcal{G}^0$ contains q'. Hence the induced hyperplane $H \cap \overline{\boldsymbol{P}'}$ for $\boldsymbol{G}'$ of $\boldsymbol{P}'$ also contains q'. Hence $\mathcal{G}^0(\boldsymbol{P}', q') = \mathcal{G}^0(\boldsymbol{P}, q)$.

Note that $\mathcal{F}^0(\boldsymbol{P}', q') = \mathcal{G}^0 \cap \boldsymbol{facets}(\boldsymbol{P}')$. The "$\subseteq$" direction is clear. No other facets have supporting hyperplanes containing q': any such facet would be in a face of $\boldsymbol{\Lambda}'$ containing q' and one of $q^{+'}$ or $q^{-'}$, in which case it would have to contain both vertices, hence be in $\mathcal{G}^0$, a contradiction. A similar argument shows that $\mathcal{R}^0(\boldsymbol{P}', q') = \mathcal{R}^0(\boldsymbol{P}, q)$.

It remains to show that any realization $\boldsymbol{P}'$ of P such that there is a point q' which conforms with $\mathcal{F}^0$ and $\mathcal{R}^0$ can be completed to a realization of $\boldsymbol{\Lambda}(\boldsymbol{P}, q)$. The idea is to follow the construction of the Lawrence extension at the beginning of this section: embed $\boldsymbol{P}'$ in $\mathbb{R}^{d+1}$ and consider two points $q^{+'}$ and $q^{-'}$ on a line which pierces the d-dimensional hyperplane $\overline{\boldsymbol{P}'}$ in q'.

We can assume that q' is a finite point: If it is not, then perform a perturbing projective transformation f on $\boldsymbol{P}'$ which makes q' finite, then construct $\boldsymbol{\Lambda}'$

and transform back by an extension of the inverse transformation f' which leaves $\boldsymbol{\Lambda}'$ finite. Such an extension is not hard to find (Lemma 4.15).

For such a finite q' Corollary 4.6 tells exactly what the facets of $\boldsymbol{\Lambda}'(\boldsymbol{P}', q')$ are. The trouble is that a priori we do not know whether $\mathcal{F}^+(\boldsymbol{P}, q) = \mathcal{F}^+(\boldsymbol{P}', q')$ and $\mathcal{F}^-(\boldsymbol{P}, q) = \mathcal{F}^-(\boldsymbol{P}', q')$. In fact they do not have to be equal like that, but it could be that $\mathcal{F}^+(\boldsymbol{P}, q) = \mathcal{F}^-(\boldsymbol{P}', q')$ and $\mathcal{F}^-(\boldsymbol{P}, q) = \mathcal{F}^+(\boldsymbol{P}', q')$. This would not constitute a problem since the face lattice would still be combinatorially equivalent with q^+ corresponding to $q^{-'}$ and q^- corresponding to $q^{+'}$. But why are these the only possibilities?

Denote $\mathcal{F}$ the set of facets of a combinatorial polytope P and $\mathcal{R}$ the set of its ridges. Each ridge is representable as the pair of the facets containing it, so by an abuse of notation $\mathcal{R} \subset \mathcal{F} \times \mathcal{F}$. The *facet graph* $\mathcal{G} = (\mathcal{F}, \mathcal{R})$ is the graph on the vertex set $\mathcal{F}$ where two facets are connected if they share a ridge. As we will show now, for a realization $\boldsymbol{P}$ of P and a point q outside of $\boldsymbol{P}$ removing $\mathcal{F}^0(\boldsymbol{P}, q)$ from the vertex set and $\mathcal{R}^0(\boldsymbol{P}, q)$ from the edge set of this graph partitions it into two connected components. The vertex sets of these components are exactly the sets $\mathcal{F}^+$ and $\mathcal{F}^-$. Once this is shown there can only be the two partitions and the proof is complete since $\mathcal{F}^+$ and $\mathcal{F}^-$ determine $\mathcal{F}^0 = \boldsymbol{facets} \setminus (\mathcal{F}^+ \cup \mathcal{F}^-)$ and $\mathcal{R}^0 = \{$all ridges contained in a facet of $\mathcal{F}^+$ and one of $\mathcal{F}^-\}$.

The facet sets $\mathcal{F}^+(\boldsymbol{P}, q)$ and $\mathcal{F}^-(\boldsymbol{P}, q)$ partition the remaining facets $\mathcal{F} \setminus \mathcal{F}^0$. Each of these sets is connected (as a subgraph): Consider a hyperplane H between $\boldsymbol{P}$ and q and project $\boldsymbol{P}$ onto this hyperplane with projection center q (see Figure 4.9). The facets in $\mathcal{F}^+$ induce a polytopal subdivision of the image of $\boldsymbol{P}$. The ridges between two facets in $\mathcal{F}^+$ are mapped to the facets of the facet images. It is easy to see that the facet graph of a polytopal subdivision is connected. Also the facets in $\mathcal{F}^-$ induce a polytopal subdivision of the image of $\boldsymbol{P}$, so also this component is connected.

We have to show that they are indeed two components, i.e. that there is no ridge connecting a facet of $\mathcal{F}^+$ and one of $\mathcal{F}^-$ which is *not* in $\mathcal{R}^0$. The defining inequalities of a facet $\boldsymbol{F}^+$ which is strictly valid for q and the defining inequality of another facet $\boldsymbol{F}^-$ which is strictly invalid for q such that the two facets share a ridge $\boldsymbol{R}$ can be convexly combined into a defining inequality of $\boldsymbol{R}$ which is satisfied with equality for q. But q cannot be in the affine hull of $\boldsymbol{R}$ or else $\boldsymbol{F}^+$ and $\boldsymbol{F}^-$ would also contain q in their affine hulls. Hence $\boldsymbol{R} \in \mathcal{R}^0$ (see Figure 4.10). $\square$

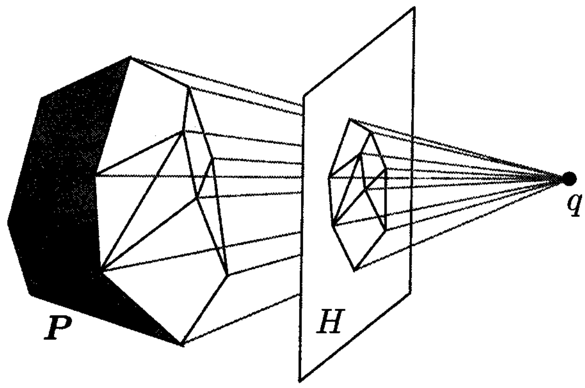

Figure 4.9: *The projection of $\boldsymbol{P}$ onto a hyperplane H induces a polytopal subdivision*

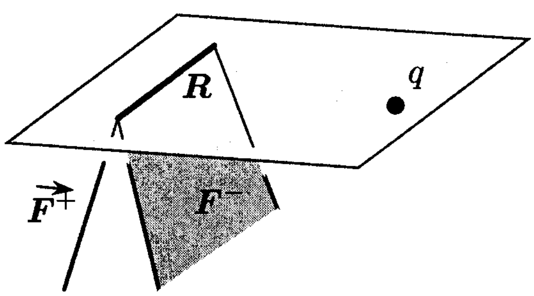

Figure 4.10: *The hyperplane defining $\boldsymbol{R}$ also containing q*

4.2.2 Connected Sum and Necessarily Flat Faces

We will show now how to use the connected sum operation in order to construct compound properties. For this we need the notion of *necessarily flat faces*.

Given two combinatorial d-polytopes P and Q whose vertex sets intersect in a common facet F, we defined the connected sum $P\#_F Q = P \cup Q \setminus \{F\}$ as the combinatorial polytope having all facets of P and Q except F (and therefore all lower-dimensional faces of both polytopes). We also say that $P\#_F Q$ is obtained by *gluing* P to Q along F.

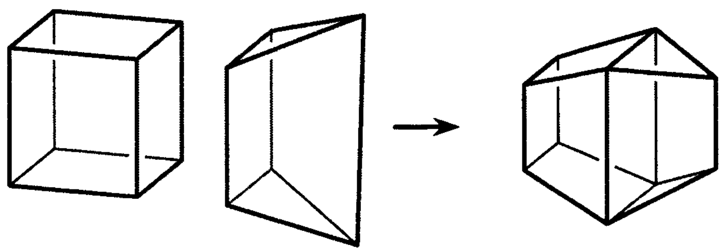

Figure 4.11: *Gluing a cube to a 3-sided prism along a square*

One way of realizing $P\#_F Q$ is to realize P and Q separately such that the realization $\boldsymbol{F}_P$ and $\boldsymbol{F}_Q$ of the common facet F are projectively equivalent. Then projective transformations can make them touch each other in F such that they lie on different sides of F and their union is convex (for instance by sending a point beyond $\boldsymbol{F}_P$ to infinity, same for $\boldsymbol{F}_Q$). See Figure 4.11.

The combinatorial structure of P may restrict the ways in which its facet F (and *its* faces) can be realized. The same holds for Q. So realizing $P\#_F Q$ in this way we have superimposed both obstructions on the faces of F from P and Q.

But are these all realizations of $P\#_F Q$? In general no, for instance the connected sum in Figure 4.11 might be perturbed in a way that the vertices on F are not coplanar (Figure 4.12). The convex hull of the vertices on F is not "flat." In this case the realization of $P\#_F Q$ cannot stem from the union of two polytopes realizing P and Q and the superimposing argument does not work.

In order for the superimposing argument to hold in *every* realization of $P\#_F Q$, the vertices of F must be realized in a $d-1$-dimensional hyperplane. This gives rise to the definition of *necessary flat* faces:

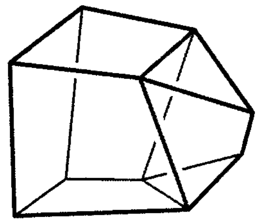

Figure 4.12: *A realization of the connected sum of cube and triangular prism such that the gluing facet is not flat — the polytopes cannot be separated*

Definition 4.7 *A d-dimensional polytope P is* necessarily flat *if for any realization of the facets of P in a space of dimension $> d$ — i.e. any realization of the vertices of F such that for each facet F of P the corresponding vertex set spans a polytope combinatorially equivalent to F — there is a d-dimensional hyperplane containing all vertices of P.*

This condition of necessary flatness is exactly what is need for the gluing facet F in order that every realization of $P \#_F Q$ can be obtained from realizing P and Q with common facet F: In any realization of $P \#_F Q$ all facets of the facet F are realized, so we can cut this polytope along the hyperplane containing all facets of F and obtain two polytopes $\boldsymbol{P}$ and $\boldsymbol{Q}$ which are combinatorially equivalent to P and Q.

Examples of Necessarily Flat Polytopes

As is easy to see, the only two-dimensional necessarily flat polytope is the triangle. However, in dimensions $d > 2$, there are well-known classes of necessarily flat polytopes: for instance pyramids, prisms, and tents. It is easy to see that the pyramid is necessarily flat, the ground face is $d-1$-dimensional, so ground face and apex together span a d-dimensional hyperplane. This argument extends to the following class of combinatorial polytopes:

Lemma 4.8 *Suppose P is a combinatorial d-polytope with two facets F and F' which contain all vertices among them and whose intersection is a ridge ($d-2$-dimensional face) R. Then P is necessarily flat.*

This class includes the tents and, for that matter, all higher-dimensional Lawrence extensions where $\mathcal{F}^0$ is non-empty.

PROOF: Suppose all of P's facets are realized in some $\mathbb{R}^n$ with $n \geq d$. Let $\boldsymbol{F}$, $\boldsymbol{F}'$, and $\boldsymbol{R}$ be the realizations of F, F', and R. Then necessarily $\boldsymbol{R} = \boldsymbol{F} \cap \boldsymbol{F}'$ and $\boldsymbol{R}$ is $d-2$-dimensional. If there is no vertex $p \in \boldsymbol{F} \setminus \overline{\boldsymbol{F}'}$ then all vertices lie in the hyperplane $\overline{\boldsymbol{F}'}$ and we are done. Otherwise we claim that the d-dimensional hyperplane $H = \boldsymbol{F}' \vee p$ contains all vertices. Note that $\overline{\boldsymbol{F}} = \boldsymbol{R} \vee p$ (the inclusion $\supseteq$ follows from $\boldsymbol{F}$ containing both $\boldsymbol{R}$ and p, also they have same dimension $d-1$ since $p \notin \boldsymbol{R}$). Since H contains both $\boldsymbol{R}$ and p, it must contain all of $\boldsymbol{F}$. The vertices of P are all in one of $\boldsymbol{F}$ and $\boldsymbol{F}'$, therefore in the d-dimensional hyperplane H. □

Lemma 4.9 *Let P be a realizable combinatorial polytope such that $\boldsymbol{prism}(P)$ has dimension $d > 2$. Then $\boldsymbol{prism}(P)$ is necessarily flat.*

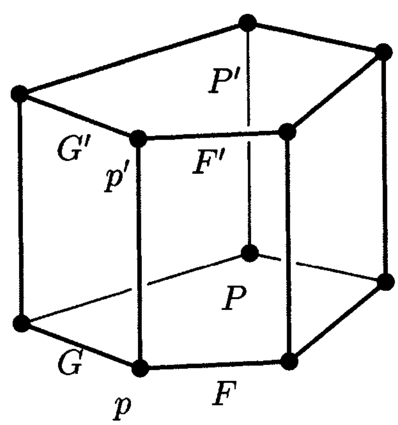

Figure 4.13: *Proof of prism being necessarily flat*

PROOF: Suppose all facets of $\boldsymbol{prism}(P)$ are realized. In particular, the realization $\boldsymbol{P}$ of the bottom face lies in a $d-1$-dimensional hyperplane $\overline{\boldsymbol{P}}$. Let p be a vertex of $\boldsymbol{P}$ and p' be the corresponding vertex in polytope $\boldsymbol{P}'$ realizing the top face. We can assume that this vertex p' is not in $\overline{\boldsymbol{P}}$: if all vertices of $\boldsymbol{P}'$ were in $\overline{\boldsymbol{P}}$ then *all* vertices would be in this $d-1$-dimensional hyperplane. We claim that the d-dimensional hyperplane $H = \boldsymbol{P} \vee p'$ contains *all* vertices. Let $\boldsymbol{F}$ be a facet of the ground face $\boldsymbol{P}$ containing p and $\boldsymbol{F}'$ the corresponding facet of the top face (which necessarily contains p'). The vertices of $\boldsymbol{F}$ and $\boldsymbol{F}'$ span a side facet of the pyramid, hence they must lie in a $d-1$-dimensional hyperplane. But this is the hyperplane spanned by the vertices of $\boldsymbol{F}$ and p', hence it is contained in H. The same argument holds for another facet $\boldsymbol{G}$ of $\boldsymbol{P}$ containing p and its corresponding facet $\boldsymbol{G}'$ of the top face. Hence we have two facets $\boldsymbol{F}'$ and $\boldsymbol{G}'$ of the top face which lie in H. Hence all of the top face must lie in H. □

There are a host of not necessarily flat polytopes. For instance all simplicial d-polytopes on more than $d+1$ vertices are not necessarily flat: In any realization of such a polytope in a d-dimensional subspace H of $\mathbb{R}^d + 1$ one point could be perturbed to be lying outside of H while the facets would still be in $d - 1$-dimensional hyperplanes – as they are simplices.

Prescribing and Gluing

Richter-Gebert in his Universality construction only glues along pyramids, prisms, and tents. Since all these are necessarily flat, the realizing polytopes can actually be split along these necessarily flat (three-dimensional) facets. So in order to get all possible realizations of the constructed polytope it is feasible to realize the building block polytopes in all possible ways with the only obstruction that the touching faces are projectively equivalent.

We want to regard the special case:

Lemma 4.10 *Let P and Q be combinatorial polytopes which share a facet F which is a pyramid over some face G. Suppose now that P and Q separately prescribe that G has some properties E_P and E_Q, respectively. Then $P\#_F Q$ prescribes that G has E_P* and E_Q.

This lemma encapsulates exactly what we meant when we said that the connected sum operation allows us to combine properties prescribed by different polytopes into one compound property. Most constructions in this thesis use this gluing along pyramids.

PROOF: Every realization of $P\#_F Q$ is the union of a realization $\boldsymbol{P}$ of P and a realization $\boldsymbol{Q}$ of Q since pyramids are necessarily flat. Hence the realization $\boldsymbol{G}$ of G must have both properties.

Let now $\boldsymbol{G}$ a realization of G with both properties and $\boldsymbol{P}$ and $\boldsymbol{Q}$ two realizations of P and Q which include $\boldsymbol{G}$ as a face. Since pyramids are projectively equivalent if and only if their ground face is projectively equivalent, we can find a projective transformation T which leaves $\boldsymbol{G}$ invariant and makes $T(\boldsymbol{P})$ and $\boldsymbol{Q}$ share the facet $\boldsymbol{F}$ (Lemma 4.15 in the appendix of this chapter). By another projective transformation leaving $\boldsymbol{F}$ invariant the polytopes can be brougth in a position where they lie on opposite sides of $\boldsymbol{F}$ and that their union is convex. □

4.3 Prescribing Edge Slopes of Polygons

We present now the link between algebra and geometry, more specifically the way polynomials equations and inequalities are encoded in polytopes.

4.3.1 Normal Polygons and Computation Frames

Richter-Gebert uses polygons to encode variables into polytopes. More precisely, he constructs combinatorial polytopes some of whose 2-faces are guaranteed to have a certain shape in every realization.

We call a polygon $G = G(1, \dots, n, 1', \dots, n')$ *normal* if the intersections $s(i) = i \wedge i'$ all lie on a common line l_∞. We call these intersection points *slopes* since if l_∞ is the line at infinity, then i and i' are parallel and $s(i)$ can be thought of as their common slope. (In this case we can identify the point $s(i)$ with homogeneous coordinates $(x, y, 0)$ with the number y/x which is exactly the common slope of the parallel lines.)

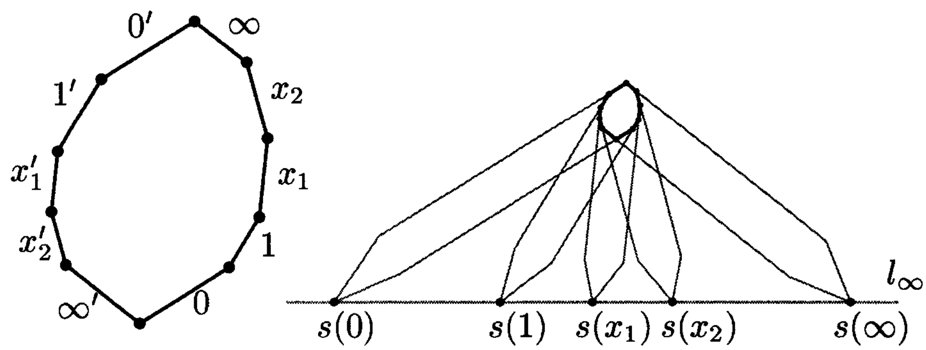

Figure 4.14: *The computation frame* $G[0, 1, x_1, \dots, x_n, \infty]$

Richter-Gebert encodes numbers in these slopes of normal polygons. All obstructions on the slopes of polygons come from the face lattice information of polytopes and therefore these slopes can only be determined up to a projectve transformation. Thus he picks three special edges whose slope act as a reference, as a *projective basis*: 0, 1, and ∞. The other edges are labeled $x_1, \dots, x_n$ (and of course there are the opposite edges $0', 1', x_1', \dots, x_n', \infty'$).

Definition 4.11 *A combinatorial polygon G is called a computation frame if its edges are labeled $0, 1, x_1, \dots, x_n, \infty, 0', 1', x_1', \dots, x_n', \infty'$, i.e. $G = G(0, 1, x_1, \dots, x_n, \infty, 0', 1', x_1', \dots, x_n', \infty')$.*

A (realized) normal polygon $\mathbf{G}$ *which is labeled as a computation frame is called* normal computation frame *(see Figure 4.14).*

In a normal computation frame $\mathbf{G}$ *let* $\alpha_i = cr(s(\infty), s(0)|s(1), s(x_i))$ *for* $i = 1, \ldots, n$ *(for the definition of the cross ratio see pages 113+). Then* $\mathbf{G}$ *is said to* encode *the assignment* $x_i = \alpha_i$.

In a very special position, if l_∞ is the line at infinity and the slopes are $s_0 = 0$ (edge 0 parallel to the x-axis), $s_1 = 1$, and $s_\infty = \infty$ (edge ∞ is parallel to the y-axis), we can just read off the variable values from these slopes: $x_i = s(x_i)$. It is easy to see that every normal computation frame can be projectively transformed to a polygon with these special slopes for 0, 1, and ∞. Hence we must have $1 < x_1 < \ldots < x_n < \infty$ (convexity argument).

Notice that if a normal computation frame encodes an assignment $x_i = \alpha_i$ then the set of points $s(0), s(1), s(x_1), \ldots, s(x_n), s(\infty)$ is projectively equivalent to the set of numbers $0, 1, \alpha_1, \ldots, \alpha_n, \infty$. This is so since $s(0)$, $s(1)$, and $s(\infty)$ constitute a projective basis and cross ratios are projectively invariant. Be aware that once these three slopes are known, an assignment $x_i = \alpha_i$ determines all other slopes.

4.3.2 Richter-Gebert's Universality Theorem

For shortness we call a system of polynomial equations and inequalities a *polynomial system.*

In his monograph [46] Richter-Gebert showed that, given a polynomial system S there is a 4-dimensional combinatorial polytope $P(S)$ such that the solution space of S is equivalent[3] to the space of all realizations of $P(S)$. He shows the theorem for polynomial systems in a special form, the *Shor normal form*:

Definition 4.12 *A polynomial system in the variables* $x_1, \ldots, x_n$ *is said to be in* Shor normal form *[46] if the only inequalities are* $1 < x_1 < \ldots < x_n < \infty$ *and if there are only equations of the form* $x_i + x_j = x_k$ *or* $x_i \cdot x_j = x_k$.

Shor [54] proved that for every polynomial system with algebraic coefficients there is a Shor normal form with a stably equivalent solution space.

For our argument we will need the following slight strengthening of the Universality Theorem.

[3]Equivalent here means *stably equivalent* which is an algebraic notion which implies homotopy equivalence, i.e. the topology is preserved [46].

Theorem 4.13 *For a system S of equations and inequalities on the variables $x_1, \ldots, x_n$ in Shor normal form, there is a combinatorial 4-polytope $P(S)$ with a 2-face $G = G[0, 1, x_1, \ldots, x_n, \infty]$ such that $P(S)$ prescribes that G is a normal computation frame encoding a solution of S.*

In other words,

1. For each realization of P the polygon realizing G is a normal computation frame and the variables values encoded in it constitute a solution of S.

2. For each solution $x_1, \ldots, x_n$ of S, each realization of G as a normal computation frame encoding this solution can be completed to a realization of $P(S)$.

The original Universality Theorem (Theorem 8.1.1 in [46]) in essence only differs in the second item:

2.' For each solution $x_1, \ldots, x_n$ *there is* a realization of G encoding this solution which can be completed to a realization of $P(S)$.

We will now examine the constructions in the proof of the Universality Theorem and see that our version of the Universality Theorem follows. These constructions are also essential ingredients in the later chapters.

Connector

We obtain the 4-dimensional *connector* polytope by gluing two 4-dimensional transmitters over n-gons (with the same n) along the prism facet. This prism facet is necessarily flat, so the connector contains the obstructions of both transmitters. The resulting polytope among others has four facets that are pyramids over the n-gon. (There are only two copies of the n-gon, but each has two pyramids incident to it.) By the transmitter property all four pyramids are projectively equivalent in every realization, but the shape of the pyramids is not prescribed further.

In our construction we will use the connector as a distributor of information: If we glue to it two polytopes along two of the pyramids then the remaining two pyramids must have the shape prescribed by both polytopes. In the polytope diagrams that we will see later we will draw the connector polytope as a dot with lines coming out of it: The lines are the pyramid facets along which we can glue more polytopes.

Gluing Diagrams

In order to visualize the gluing in four dimensions when there are many partaking polytopes we use gluing diagrams. There polytopes are just boxes which are connected by lines. In our constructions in Chapters 5 and 6 we only use pyramids as glue. So the lines between the polytope boxes always mean pyramids over polygons. For the connector polytope we have four pyramids, its symbol is just a dot. Figure 4.15 shows the gluing diagram of four polytopes which all have a pyramid face over an n-gon G which each of them prescribes to have a certain property. By gluing them together G is now prescribed to have *all* four properties.

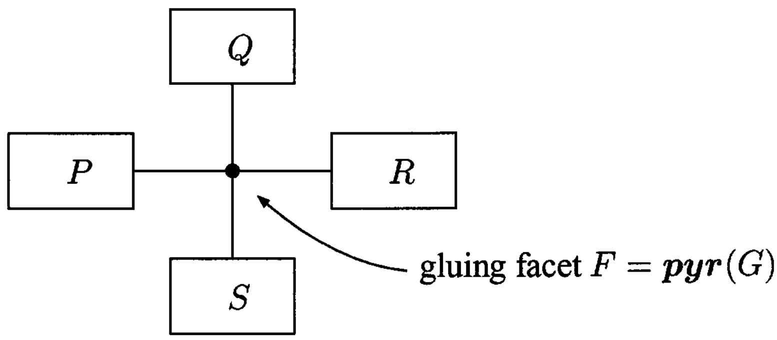

Figure 4.15: *Gluing diagram which prescribes many properties on G*

Edge Forgetter

By a construction similar to the transmitter we obtain the *edge forgetter*. (Richter-Gebert calls this polytope the *forgetful transmitter*; we renamed it because we contrast it in later chapters with the *vertex forgetter*.) This polytope is constructed as follows: Take an upright prism over a polygon $G(1, 2, \ldots, i-1, i+1, \ldots, n)$. Now cut off the vertex $i-1 \wedge i+1$. The new facets introduces a new edge to the bottom polygon, say i. The edge forgetter $E \setminus i$ is the Lawrence extension of this mutilated prism where q is the intersection of the edges connecting the top and bottom faces (and the one that was cut off). Notice that for this Lawrence extension $\mathcal{F}^0$ contains the side facets of the prism (without the facet introduced by the cut-off vertex), $\mathcal{F}^+$ are the bottom polygon and the new triangular facet, and $\mathcal{F}^-$ is the top polygon. The ridge set $\mathcal{R}^0$ is empty. See Figure 4.16.

By a Lawrence extension argument much like the one about the transmitter

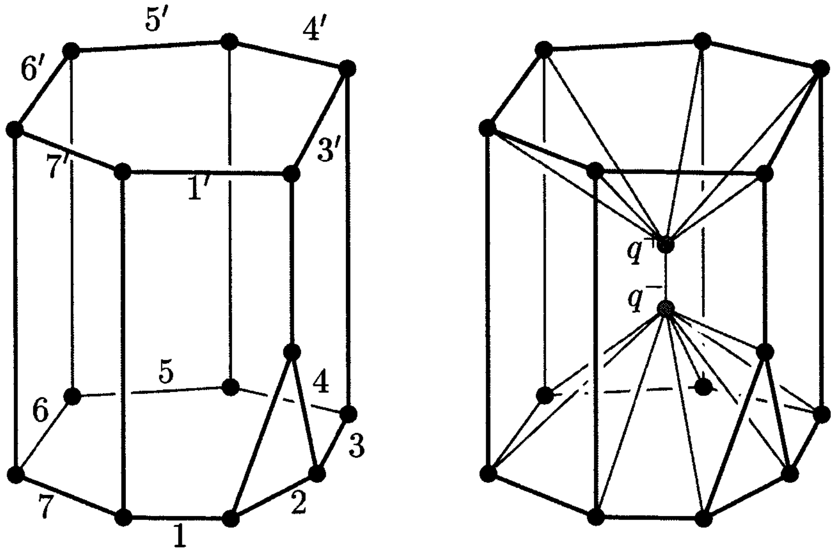

Figure 4.16: *Edge forgetter* $E \setminus 2$

this polytope prescribes for the mutilated prism facet that the lines supporting the edges in the top polygon are *projectively equivalent* to their corresponding lines in the bottom polygon. Of course, the bottom polygon has one edge more, this edge has been *forgotten*.

We can glue two of these edge forgetters on top of each other along a common pyramid. By stacking more edge forgetters on top of each other many edges $i_1, \ldots, i_r$ can be forgotten. The resulting polytope $E \setminus i_1, \ldots, i_r$ prescribes that the the edges of the polygon on the small end are projectively equivalent to the corresponding edges of the polygon on the large end.

In the polytope diagrams we will draw the forgetters as trapezoids: the small end signifies the pyramid over the polygon where edges or vertices have been forgotten, the large end signifies the pyramid with a complete edge or vertex set (see Figure 4.17).

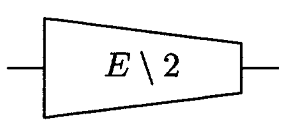

Figure 4.17: *Symbol for the edge forgetter*

Adding and Multiplying

The most involved part in Richter-Gebert's construction considers the polytopes that encode addition and multiplication: For addition he constructs the combinatorial 4-polytope $P^{x+y}[0, 1, x_1, x_2, x_3, \infty]$ which contains a pyramid facet over a 12-gon $G = G[0, 1, x_1, x_2, x_3, \infty]$. This polytope prescribes that G is a normal computation frame encoding that the variables x_1, x_2, and x_3 satisfy $x_1 + x_2 = x_3$ (see Theorem 7.1.1 in [46]). Also he constructs the combinatorial 4-polytope and $P^{x \cdot y}[0, 1, x_1, x_2, x_3, \infty]$ which contains a pyramid facet over a 12-gon $G = G[0, 1, x_1, x_2, x_3, \infty]$ and prescribes that G is normal and encodes that $x_1 \cdot x_2 = x_3$ (see Theorem 7.2.2 in [46]). Figure 4.18 shows the diagram pieces for these polytopes, the lines going out of the boxes are of course the pyramids over the 12-gons.

$$\boxed{P^{x+y}[0, 1, x_1, x_3, x_3, \infty]} \text{---} \qquad \boxed{P^{x \cdot y}[0, 1, x_1, x_3, x_3, \infty]} \text{---}$$

Figure 4.18: *Symbols for adding and multiplying polytopes*

How to encode a concrete Shor Normal form

We now give an example which showcases the proof technique of Theorem 4.13 and should make the general construction evident. Figure 4.19 shows the schematic view of the constructed polytope for the Shor normal form

$$\begin{gathered} 1 < x_1 < x_2 < x_3 < x_4 < x_5 < \infty, \\ x_1 + x_3 = x_4, \\ x_2 + x_4 = x_5, \\ x_2 \cdot x_3 = x_5 \end{gathered}$$

From the connectors on the right-hand side that share the computation frame $G[0, 1, x_1, x_2, x_3, x_4, x_5, \infty]$ we have three edge forgetters, each of which is glued to an adding or multiplying polytope. Notice first that this polytope prescribes that the computation frames of the connectors be normal: The points $0 \wedge 0'$, $1 \wedge 1'$, $x_1 \wedge x_1'$, $x_3 \wedge x_3'$, $x_4 \wedge x_4'$, and $\infty \wedge \infty'$ are on a line by the adder $P^{x+y}[0, 1, x_1, x_3, x_4, \infty]$, the other points $x_2 \wedge x_2'$ and $x_5 \wedge x_5'$ are on the same line spanned by $0 \wedge 0'$ and $1 \wedge 1'$ by the other adder. Also the polytope prescribes that *all three equations* in the computation frames of the connectors are prescribed to be valid. By the computation frame encoding trivially $1 < x_1 < \ldots < x_n < \infty$.

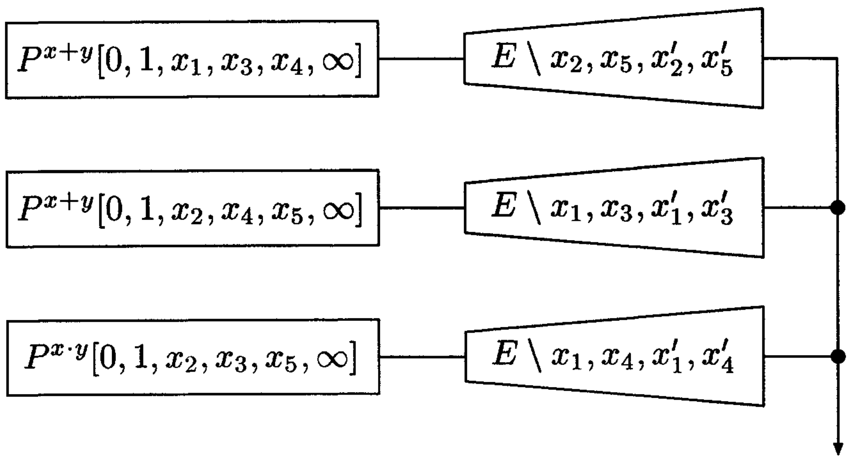

Figure 4.19: *Polytope prescribing the solution for a concrete Shor normal form*

4.4 Appendix: Projective Space and Polytopes

Realizations of combinatorial polytopes can never be unique: a translation or a rotation of a polytope $\boldsymbol{P}$ which realizes a combinatorial polytope P results in other realizations of P. Affine transformations of $\mathbb{R}^d$ (linear transformations followed by a translation) preserve the face lattice of a polytope. Even more general transformation that leave the face lattice of polytope invariant are *projective transformations* which act on the projective space $\mathbb{R}P^d$.

We first give an introduction to projective spaces, projective subspaces and the projective transformations between them. We will be most interested in the various ways of constructing projective transformations for its use in later constructions. Also we briefly explain cross ratios and quadrangular sets. Finally, we show how to embed convex polytopes in the projective setup.

4.4.1 Projective Space

We want to give a short introduction to projective spaces and projective transformations. We assume the reader is familiar with affine geometry (points, lines, planes not necessarily through the origin of some $\mathbb{R}^d$). The points in $\mathbb{R}^d$ are called finite points. It is often helpful, however, to include *infinite points*. These can be thought of as equivalence classes of lines where two lines are equivalent if they are parallel. When we say "parallel lines meet at infinity" we mean that for each *direction* of parallel lines we add *one point at infinity*.

The union of the finite and infinite points is the d-dimensional *projective space* $\mathbb{R}P^d$.

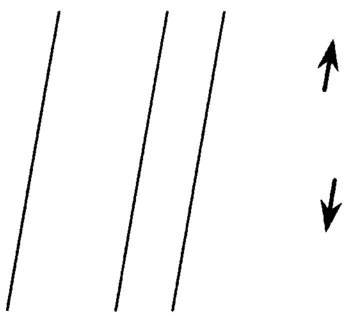

Figure 4.20: *Construction of $\mathbb{R}P^2$ from $\mathbb{R}^2$ by adding points at infinity*

The d-dimensional projective space also be obtained as the set of all 1-dimensional linear subspaces (lines through the origin) of the space $\mathbb{R}^{d+1}$ (or any $d+1$-dimensional vector space over $\mathbb{R}$) — or the quotient space of $\mathbb{R}^{d+1}$ where scalar multiples of vectors are identified.

A usual translation between the two approaches is *homogenization*: Embed the d-dimensional affine space $\mathbb{R}^d$ in the $z_{d+1} = 1$ plane of the vector space $\mathbb{R}^{d+1}$. Each *finite* point in $\mathbb{R}P^d$ corresponds to a certain line through the origin and the lifted point in $\mathbb{R}^{d+1}$, namely

$$(x_1, \ldots, x_d) \in \mathbb{R}P^d \quad \rightsquigarrow \quad \mathbb{R}(x_1, \ldots, x_d, 1) \subset \mathbb{R}^{d+1}$$

Hence the *finite* points are the ones having representing $d+1$-vectors with non-zero last coordinate. On the other hand, a point at infinity corresponds to a line containing only vectors with last coordinate zero. By a projective point x we will from now on mean a set of the form $\mathbb{R}(x_1, \ldots, x_{d+1})$ where of course the $x_1, \ldots, x_{d+1}$ are only determined up to common scalar multiple.

4.4.2 Projective Subspaces

The linear subspaces of $\mathbb{R}^{d+1}$ induce the *projective subspaces* (*flats*) of $\mathbb{R}P^d$. This is possible since linear subspaces are closed under scalar multiples. For a subset S of $\mathbb{R}P^d$ we denote by $\overline{S}$ the *projective closure* of S which is defined as the smallest flat that contains S. The subsets of these projective subspaces consisting of their *finite* points are then the *affine subspaces* of $\mathbb{R}^d$: points, lines, planes, hyperplanes etc. The smallest affine subspace containing a set S

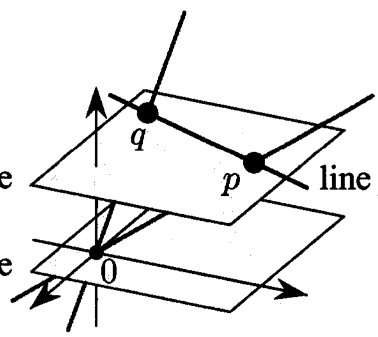

Figure 4.21: *Homogenization*

of finite points is called the *affine hull* of S. We will not distinguish between the affine and projective subspaces as well as between the affine hull and the projective closure. By this token we will use these names of affine subspaces interchangeably with their projective closures. Now it becomes clear why it is sensible to talk about the hyperplane "at infinity", it is induced by all vectors in $\mathbb{R}^{d+1}$ having last coordinate zero which is a linear subspace, more precisely a $d-1$-dimensional hyperplane. Projective subspaces are again projective spaces, namely if a flat V_{proj} is induced by a linear subspace V_{lin} of $\mathbb{R}^{d+1}$, then V_{proj} can be viewed as the set of 1-dimensional subspaces of V_{lin}, i.e. $V_{proj} \equiv V_{lin}/\mathbb{R}$. Hence, whenever we talk about the projective space $\mathbb{R}^d$, this space acts as a representative for any d-dimensional flat.

The union of projective subspaces in general is not a projective subspace. We call the *span* or *join* of flats V and W of $\mathbb{R}P^d$ the smallest projective subspace containing all points of $V \cup W$. Denote the join of V and W by $V \vee W$. If S and T are mere subsets of $\mathbb{R}P^d$, we abbreviate $\overline{S} \vee \overline{T}$ by $S \vee T$. (See Figure 4.21 for the join of two points.) On the other hand, the intersection or *meet* of projective subspaces is always a projective subspace. Denote the meet of two flats V and W by $V \wedge W$. Again for mere subsets S and T of $\mathbb{R}^d$ abbreviate $\overline{S} \wedge \overline{T}$ by $S \wedge T$.

We say a set of points in $\mathbb{R}P^d$ is in *general position* if the span of any d-element subset contains no other of the remaining points. This is equivalent to the fact that the representing vectors in $\mathbb{R}^{d+1}$ are also in general position, i.e. that every $d+1$-element subset of them is a linear basis.

4.4.3 Projective Transformations

Linear transformation between linear subspaces of $\mathbb{R}^{d+1}$ commute with multiplying vectors with real scalars. Hence they induce transformations of the projective subspaces, the *projective transformations*. It is a fundamental theorem they are *exactly* those transformations that take projective subspaces to projective subspaces. It is this property that also implies that the meet and projective transformations as well as join and projective transformations commute, i.e. for subsets S and T and a projective transformation f we have $f(S \wedge T) = f(S) \wedge f(T)$ and $f(S \vee T) = f(S) \vee f(T)$. Naturally, for a projective transformation $f : \mathbb{R}P^d \to W$ the restriction $f|_V : V \to W$ for a subspace V of $\mathbb{R}P^d$ is a projective transformation. We call a projective transformation *invertible* if the inducing linear transformation is invertible. This is equivalent to the fact that the projective transformation preserves the dimensions of projective subspaces. Most of the projective transformations we will use will be invertible (we will draw the reader's attention to the rare cases where they are not). We say that two sets S and T are *projectively equivalent* if there is an invertible projective transformation f such that $f(S) = T$.

The $d-1$-dimensional hyperplanes of $\mathbb{R}P^d$ are in a natural one-to-one correspondence with the points of $\mathbb{R}P^d$: They are induced by a d-dimensional subspace of $\mathbb{R}^{d+1}$ which can be represented by their orthogonal 1-dimensional subspaces. By this token, a point $p \in \mathbb{R}P^d$ represented by a vector $(p_1, \ldots, p_{d+1})$ is on a $d-1$-dimensional hyperplane H represented by a vector $(a_1, \ldots, a_{d+1})$ if and only if these vectors are orthogonal: $a_1p_1 + \ldots + a_{d+1}p_{d+1} = 0$. This duality of points and $d-1$-dimensional hyperplanes carries over many results about points in projective space to corresponding results about $d-1$-dimensional hyperplanes. A projective transformation corresponding to a linear transformation with invertible matrix M takes a hyperplane represented by a vector $a = (a_1, \ldots, a_{d+1})$ to the hyperplane with vector $M^{-1}a$.

The next three lemmas are tools to construct projective transformations. The first lemma tells us that the image of $d+2$ points in general position in a projective space of dimension d uniquely determines the projective transformation.

Lemma 4.14 *Let $v_1, \ldots, v_{d+2} \in \mathbb{R}P^d$ be in general position and $w_1, \ldots, w_{d+2} \in \mathbb{R}P^d$ in any position. Then there is exactly one projective transformation f such that $f(v_i) = w_i$ for all $i = 1, \ldots, d+2$.*

If the w_i are in general position then f is invertible.

This lemma gives rise to the concept of a *projective basis* of a d-dimensional projective space: it is a set of $d+2$ points in general position. Whenever a configuration of points in $\mathbb{R}P^d$ is determined up to a projective transformation it is enough to fix the position of $d+2$ points in general position and then the position of all other points is fixed. For instance, in one dimension it is enough to fix the position of three points, the position of all other points is then determined by their cross ratio with the projective basis. Note that this lemma also holds in the dualized version, i.e. when the points are substituted by $d-1$-dimensional hyperplanes.

PROOF: We must find a linear transformation f such that for any representing vectors $v_1, \ldots, v_{d+2}$ and $w_1, \ldots, w_{d+2}$ of the projective points with the same symbols we have $f(v_i) = \lambda w_i$.

By general position the representing vectors $v_1, \ldots, v_{d+1}$ are independent and linearly span $\mathbb{R}^{d+1}$. Hence there are scalars $\mu_1, \ldots, \mu_{d+1}$ such that $v_{d+2} = \mu_1 v_1 + \ldots + \mu_{d+1} v_{d+1}$. By another general position argument the μ_i are unique ($v_1, \ldots, v_{d+1}$ are linear independent) and nonzero (for all i the vectors $v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_{d+1}, v_{d+2}$ are independent, hence μ_i cannot be zero). Also there are unique and nonzero scalars $\nu_1, \ldots, \nu_{d+1}$ such that $w_{d+2} = \nu_1 w_1 + \ldots + \nu_{d+1} w_{d+1}$.

There is exactly one linear transformation mapping the vectors $\mu_i v_i$ to the vectors $\nu_i w_i$ for $i = 1, \ldots, d+1$. By linearity this transformation also takes v_{d+2} to w_{d+2}. Uniqueness is also not hard to see: If a linear transformation took the vectors $\mu_i v_i$ to vectors different from a common scalar multiple of $\nu_i w_i$ then it would not take v_{d+2} to a scalar multiple of w_{d+2}. Hence the linear transformation is unique up to common scalar multiple. □

The second lemma tells us how we can extend a projective transformation defined on a subspace V to a projective transformation mapping an additional point $p \notin V$ to an arbitrary vertex q.

Lemma 4.15 *Let V be a projective subspace of $\mathbb{R}P^d$, f be an invertible projective transformation on V, p and q some points outside V and p' and q' some points outside of $f(V)$ such that $f(V \wedge (p \vee q)) = f(V) \wedge (p' \vee q')$. Then there is a unique projective transformation g on $V \vee p$ such that $g|_V = f$, $g(p) = p'$, and $g(q) = q'$.*

PROOF: Let d_V be the dimension of V. Let $v_0 = V \wedge (p \vee q)$ and $w_0 = f(V) \wedge (p' \vee q')$. Pick $d_V + 1$ points $v_1, \ldots, v_{d_V+1}$ such that the $d_V + 2$ points $v_0, v_1, \ldots, v_{d_V+1}$ in V are in general position. Then the points

$p, q, v_1, \ldots, v_{d_V+1}$ are in general position and span $V \vee p$. Similarly the images $w_i = f(v_i)$ for $i = 1, \ldots, d_V + 2$ and the points p' and q' are in general position. See Figure 4.22.

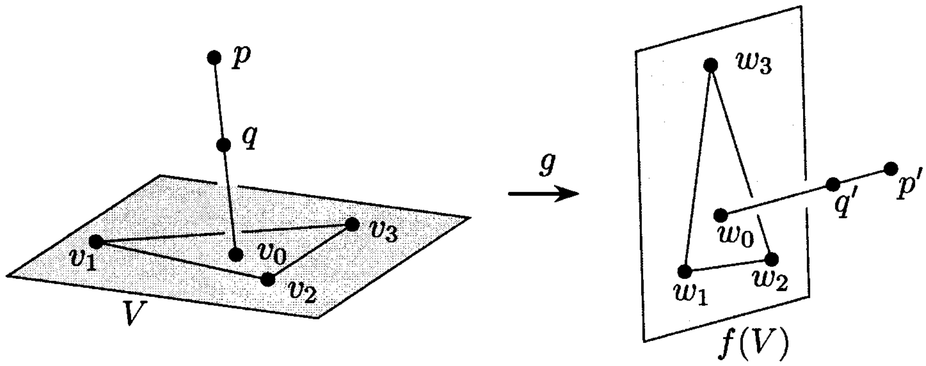

Figure 4.22: *Construction of g in the Proof of Lemma 4.15*

By the previous lemma there is a projective transformation g such that $g(p) = p', g(q) = q', g(v_1) = w_1, \ldots, g(v_{d+1}) = w_{d+1}$. Note that g maps V to $f(V)$ and the line $p \vee q$ to the line $p' \vee q'$ (join and projective transformations commute). Hence it also maps the meet $v_0 = V \wedge (p \vee q)$ to $w_0 = f(V) \wedge (p' \vee q')$. Since again by the previous lemma there is only one projective transformation in V taking $v_0, v_1, \ldots, v_{d_V+1}$ to $w_0, w_1, \ldots, w_{d_V+1}$, the transformation g must equal f on V. □

As an easy corollary we note that if the image of only one point outside V is known then there are many projective transformations extending f:

Corollary 4.16 *Let V be a projective subspace of $\mathbb{R}P^d$, f be an invertible projective transformation on V, p some point outside V and p' some point outside of $f(V)$. Then there is a projective transformation g on $V \vee p$ such that $g|_V = f$ and $g(p) = p'$.*

The last of the lemmas in this section will be crucial in the sequel: In $\mathbb{R}P^3$, under some assumptions a projective transformation is uniquely determined by the images of four points in general position and the image of a plane.

Lemma 4.17 *Let $v_1, \ldots, v_{d+1} \in \mathbb{R}P^d$ in general position and H be a hyperplane that misses all of them. Furthermore, let $w_1, \ldots, w_{d+1} \in \mathbb{R}P^d$ also be in general position and G be a hyperplane missing all of them.*

Then there is a unique projective transformation f with $f(v_i) = w_i$ for all $i = 1, \ldots, d+1$ and with $f(H) = G$.

PROOF: Let H_i be the hyperplane spanned by the vertices except v_i: $H_i = \bigvee_{k\neq i} v_k$. The hyperplanes $H_1, \ldots, H_n, H$ are in general position. This is so since no meet on d of these hyperplanes is contained in the two remaining hyperplanes: (1) $\bigwedge_{k\neq i} H_k = v_i$ is not contained in either H_i (the v_j are in general position) or H (assumption that $v_i \notin H$) and (2) $\bigwedge_{k\neq i,j} H_k \wedge H = (v_i \vee v_j) \wedge H$ is not on H_i or else v_j lies on H (H_j analogously).

Define similarly $G_i = w_1 \vee \ldots \vee w_{i-1} \vee w_{i+1} \vee \ldots \vee w_{d+1}$, then $G, G_1, \ldots, G_{d+1}$ are also in general position. If there is a projective transformation with $T(v_i) = w_i$, then by

$$\begin{aligned} T(H_i) &= T(v_1 \vee \ldots \vee v_{i-1} \vee v_{i+1} \vee \ldots \vee v_{d+1}) \\ &= w_1 \vee \ldots \vee w_{i-1} \vee w_{i+1} \vee \ldots \vee w_{d+1} \\ &= G_i \end{aligned}$$

it has to map H_i to G_i. By the dualized version of Lemma 4.14 there is exactly one projective transformation T with $T(H) = T(G)$ and $T(H_i) = G_i$. Note that then necessarily

$$\begin{aligned} T(v_i) &= T(H_1 \wedge \ldots \wedge H_{i-1} \wedge H_{i+1} \wedge \ldots \wedge H_{d+1}) \\ &= G_1 \wedge \ldots \wedge G_{i-1} \wedge G_{i+1} \wedge \ldots \wedge G_{d+1} \\ &= w_i. \end{aligned}$$

□

We should also talk about very special projective transformations, the *projections*.

Lemma 4.18 *Let V and V' two $d-1$-dimensional projective subspaces of $\mathbb{R}P^d$ and c a point outside of V and V'. Then the function $f : V \to V'$, $f(p) = (p \vee c) \wedge V'$ is an invertible projective transformation.*

PROOF: Let W be a $d-1$ dimensional projective subspace which intersects both V and V' in $V \cap V'$, but which does not contain c. By Lemma 4.15 there is a unique projective transformation $g : \mathbb{R}P^d \to \mathbb{R}P^d$ which extends the identity map of W and map c back to c and some point p in V to $f(p)$. This projective transformation takes V to V' since $V = (V \wedge W) \vee p$ and $g(V) = g(V \wedge W) \vee g(p) = (V \wedge W) \vee f(p) = (V' \wedge W) \vee f(p) = V'$. We claim that $g|_V = f$: Let $q \in V$. Then $q_W := (q \vee c) \wedge W$ lies in W, hence $g(q_W) = id(q_W) = q_W$. But $q = (q_W \vee c) \wedge V$, hence $g(q) = (g(q_W) \vee g(c)) \wedge g(V') = (q \vee c) \wedge V' = f(q)$. □

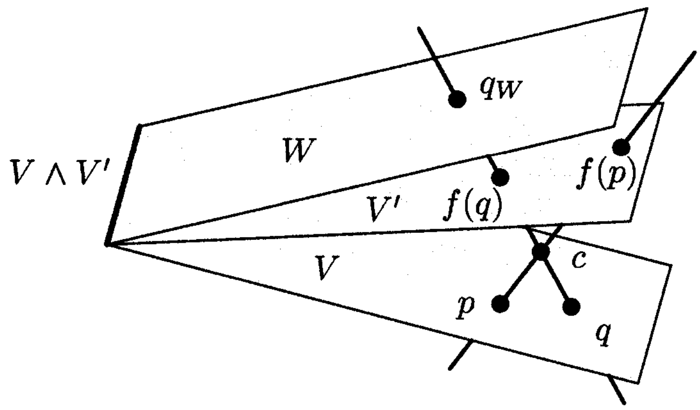

Figure 4.23: *Proof of Lemma 4.18*

4.4.4 Cross Ratios and Quadrangular Sets

For four points $p_1, p_2, p_3, p_4 \in \mathbb{R}P^1$ with representative vectors p_1^{rep}, p_2^{rep}, $p_3^{rep}, p_3^{rep} \in \mathbb{R}^2$ the *cross ratio* is defined as

$$cr(p_1, p_2|p_3, p_4) = \frac{\det(p_1^{rep}, p_3^{rep}) \cdot \det(p_2^{rep}, p_4^{rep})}{\det(p_1^{rep}, p_4^{rep}) \cdot \det(p_2^{rep}, p_3^{rep})}.$$

This definition makes only sense since different representatives differ only by nonzero scalar multiples which cancel out in the quotient. If the p_i are finite points the cross ratio can be computed by using only the signed (!) euclidean distances of the points:

$$cr(p_1, p_2|p_3, p_4) = \frac{|p_1, p_3| \cdot |p_2, p_4|}{|p_1, p_4| \cdot |p_2, p_3|}$$

since $\det(\binom{p_1}{1}, \binom{p_3}{1}) = p_3 - p_1 = |p_1, p_3|$. Cross ratios are evidently invariant under invertible projective transformations of $\mathbb{R}^1$. Using this, we can define the cross ratio of four points on a line in a higher-dimensional projective space by projecting the line down to $\mathbb{R}P^1$ and computing the cross ratio there. Note that these general cross ratios are invariant under invertible projective transformations of the line containing the four points.

Let p, q, r, s four points on a line such that p, q, and r are distinct. Then the cross ratio $cr(p, q|r, s)$ determines the position of s: The unique projective transformation f taking p to $\infty \in \mathbb{R}P^1$, q to 0 and r to 1 leaves the cross ratio invariant, hence $cr(p, q|r, s) = cr(\infty, 0|1, f(s)) = f(s)$. So $f^{-1}(cr(p, q|r, s))$ determines s.

Finally, we want to briefly talk about quadrangular sets: Six points on a line l are called a *quadrangular set* if they are the intersections of l with the six possible lines formed by four distinct points. We note two properties of quadrangular sets:

Lemma 4.19 *1. Five points in a quadrangular set determine the sixth point.*

2. Quadrangular sets are invariant under projective transformations.

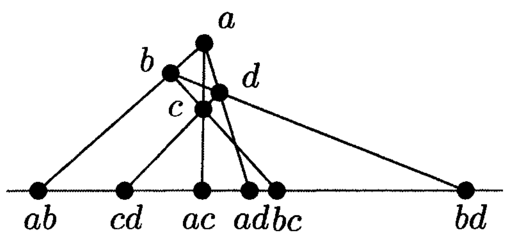

Figure 4.24: *The points ab, cd, ac, ad, bc, bd are a quadrangular set*

PROOF: If the points are labeled as in Figure 4.24 then we can define two projections, one uses a as projection center and projects l to $b \vee d$ and the other uses c as projection center and projects $b \vee d$ back to l. If we label $p = (a \vee c) \wedge (b \vee d)$ then since projections leave cross ratios invariant we have $cr(bd, ab|ac, ad) = cr(bd, b|p, d) = cr(bd, bc|ac, cd)$. Note now that if all points in the quadrangular set except for instance cd are known, then also these cross ratios are known and therefore also cd. By symmetry this holds for all other points as well. It also follows that quadrangular sets are invariant under projective transformations. □

4.4.5 Polytopes in Projective Space

A polytope is equivalently the convex hull of a finite number of (finite) points and the intersection of a finite number of halfspaces such that the result is bounded. How does this fit in with the projective closure of the affine real space? A hyperplane does not cut the projective space into two components, it stays connected. This is where the hyperplane at infinity comes into play: A hyperplane does cut the set of *finite* points into two two components. We call these components *halfspaces*. The convex hull is similarly defined: two finite points p and q cut the line $p \vee q$ into two connected pieces, we call the piece which does not contain an infinite point the *line segment between* p and q . The following lemma gives a criterion for betweenness.

Lemma 4.20 *If* p, q, r, s *are distinct points on a line such that* p *and* q *are finite and* r *is not between* p *and* q. *Then* s *is not between* p *and* q *if and only if* $cr(p, q|r, s) > 0$.

The lemma is an easy implication on the formula of the cross ratio using signed euclidean distances.

Projective transformations do not in general map polytopes back to polytopes: points of the polytope might get mapped to infinity and the result would not be convex. We define *admissible* projective transformation of a polytope $\boldsymbol{P}$ according to Ziegler [61] as projective transformations that do map $\boldsymbol{P}$ back to a polytope. We note the following lemma from Ziegler's treatment:

Lemma 4.21 *Let* $\boldsymbol{P}$ *be a polytope in* $\mathbb{R}^d \subset \mathbb{R}P^d$ *and* f *a projective transformation such that the preimage of the hyperplane at infinity under* f *does not intersect* $\boldsymbol{P}$. *Then the image* $f(\boldsymbol{P})$ *consist only of finite points and* $f(\boldsymbol{P})$ *is the convex polytope spanned by the images of the vertices of* $\boldsymbol{P}$, *i.e.* $f(\boldsymbol{P}) = conv(f(vert(\boldsymbol{P})))$.

If two polytopes $\boldsymbol{P}$ and $\boldsymbol{P}'$ are projectively equivalent, the invertible projective transformation f with $f(\boldsymbol{P}) = \boldsymbol{P}'$ induces a one-to-one correspondence between the faces of $\boldsymbol{P}$ and the faces of $\boldsymbol{P}'$, i.e. for all faces F of $\boldsymbol{P}$ the image $f(F)$ is a face of $\boldsymbol{P}'$ with the same dimension as F and vice versa.

Chapter 5

Prescribing Faces of Polytopes

In this chapter we prove the following two results concerning the prescribability of the *exact shape* of polytopes:

1. Let $\boldsymbol{G}$ be a (realized) polygon with algebraic vertex coordinates. Then we can construct a combinatorial 4-dimensional polytope $P(G)$ which has a face G which it prescribes to be projectively equivalent to $\boldsymbol{G}$ (Section 5.2).

2. Let $d \geq 3$. Let $\boldsymbol{G}$ be a d-dimensional (realized) polytope with algebraic vertex coordinates. Then we can construct a combinatorial $d+2$-dimensional polytope $P(G)$ which has a face G which it prescribes to be projectively equivalent to $\boldsymbol{G}$ (Section 5.3).

Note that these results are best possible in two ways: First only projective properties can be prescribed by combinatorial polytopes, so we cannot ask more than prescribing $\boldsymbol{G}$ up to projective equivalence. Second a combinatorial polytope cannot prescribe one of its faces to be projective to a polytope with non-algebraic coordinates (more specifically a polytope every projective equivalence of which has non-algebraic coordinates): In the next section we will show how the decision algorithm for the existential theory of the reals by Tarski (and later Seidenberg) [13] implies that if a polytope has a realization with non-algebraic coordinates then it must also have a realization with purely algebraic coordinates.

By a variant of Steinitz' theorem [3] it is impossible to prescribe 2-faces of 3-polytopes. It remains unknown, however, whether *facets* of 4-polytopes or higher-dimensional polytopes can be prescribed. Some polytopes occuring as these facets which *can* be prescribed are prisms and the "taco" polytopes depicted in Figure 5.9. No other classes are known to us.

5.1 Algebraic Numbers and the Existential Theory of the Reals

Real algebraic numbers are real roots of rational univariate polynomials. Since we are only concerned with real numbers here (or more precisely only with ordered fields) we will just say *algebraic numbers* when we mean real algebraic numbers. Each real algebraic number α can be separated by an irreducible rational polynomial f and two bounding rational numbers, i.e. α is represented by a triple (f, l, r) such that α is the *only* number with $f(\alpha) = 0$ and $l < \alpha < r$.

The *existential (problem of the first-order) theory of the reals* is the decision problem whether a set of polynomial equations ($f(x_1, \ldots, x_n) = 0$) and polynomial inequalities ($f(x_1, \ldots, x_n) < 0$ or ≤ 0) with coefficients in a real-closed field has a solution in this field. (Real-closed fields are ordered fields where every positive number has a root and where every odd-degree polynomial has a root.) The Tarski-Seidenberg theorem [53] gives a finite decision procedure for this problem which only uses the axioms of the real-closed fields. Since both the field of real algebraic numbers as well as the field of real numbers are real-closed, a polynomial system either has a solution in the algebraic numbers or no solution at all (not even in the real numbers). A nice survey for this subject is [40].

Theorem 5.1 *A combinatorial polytope either has a realization using only real algebraic coordinates or it has no realization at all, not even in the real numbers.*

PROOF: The realization space of a combinatorial polytope is the solution set of polynomial equations and inequalities: For each $d+1$-tuple $(v_1, \ldots, v_d, v_{d+1})$ of vertices lying on a common facet we get the equation

$$\det \begin{pmatrix} v_1 & \cdots & v_d & v_{d+1} \\ 1 & \cdots & 1 & 1 \end{pmatrix} = 0$$

where $v_i \in \mathbb{R}^d$. For each d-tuple of vertices on a facet and a $d+1$st vertex v_{d+1} not in this facet the equation is changed into a < 0. (The d-tuple must be oriented correctly, but it is easy to find a consistent orientation for each d-tuple of vertices in a facet from the realized polytopes of which the combinatorial polytope is composed.)

Now by the Tarski-Seidenberg theorem this polynomial system either has a real algebraic solution – in which case it is realizable over the algebraic numbers – or no real solution at all – in which case it is nonrealizable, even over the real numbers. □

5.2 4-Polytopes Prescribing 2-Faces

In this section we will show the prescribability of polygons in 4-polytopes. We will give an overview of the proof now:

The first goal will be, given algebraic numbers $1 < \alpha_1 < \ldots < \alpha_m$, to construct a 4-polytope $P(\alpha_1, \ldots, \alpha_m)$ which prescribes one of its polygonal faces to be normal and have edge slopes which are projectively equivalent to $0, 1, \alpha_1, \ldots, \alpha_m, \infty$. (The edge slopes are points on a line, the α_i are points on the real line, so it makes sense to speak of projective equivalence.) We do this by constructing a polytope which prescribes one of its faces to be a normal computation frame that encodes the assignment $x_i = \alpha_i$ for $i = 1, \ldots, m$.

In Section 5.2.1 we will show a lemma which implies that for a polynomial system that defines the algebraic numbers $\alpha_1 < \ldots < \alpha_m$ there is a Shor normal form $S = S(\alpha_1, \ldots, \alpha_m)$ in variables $x_1 < \ldots < x_n$ (with $n > m$) such that the only solution for $x_1, \ldots, x_m$ are the algebraic numbers $\alpha_1, \ldots, \alpha_m$ (Corollary 5.3).

For this Shor normal form we can invoke Richter-Gebert's universality theorem (Theorem 4.13): In Section 5.2.2 we will use edge forgetters to forget all edges corresponding to auxiliary variables (variables x_i with $i > m$), this will give us a polytope which prescribes a polygon to be normal and to have the exact edge slopes $0, 1, \alpha_1, \ldots, \alpha_m, \infty$ (up to projective equivalence).

In order to make the leap from prescribing the slopes of the edges to prescribing the actual edges we use a new kind of gadget, the *vertex forgetter*. We introduce it in Section 5.2.3 and put it to work in Section 5.2.4. In the latter section we talk about prescribing a *centrally symmetric* polygon $G^{\diamond}$. This polygon is of course normal — the line containing the edge slopes is at infinity. But also each diagonal connecting two vertices v and w is parallel to the diagonal connecting the vertices opposite to v and w, hence these diagonals also meet on this line. We call these intersection points the *chord slopes*. If we delete a vertex and its opposite from the set of vertices and take the convex hull we get a polygon which has a 1-chord of the original polytope on its boundary and therefore still is a normal polygon.

We can presribe the edge slopes of this polytope (if the vertex coordinates are algebraic numbers then all edge slopes and chord slopes are also algebraic). To this polytope we glue a new kind of polytope the *vertex forgetter*: this new kind of polytope works a bit like the edge forgetter, only it forgets vertices rather than edges, and all other vertices are projectively equivalent.

By gluing this to a polytope which prescribed the original edge slopes of $G^{\Diamond}$ we obtain a polytope which prescribes the edge slopes *and a* 1*-chord slope*. We use connectors and many polytopes which presribe different 1-chord slopes and obtain one big polytope $P(G^{\Diamond})$ which prescribes that one of its faces has all edge slopes and 1-chord slopes of $G^{\Diamond}$ (up to projective equivalence).

In Section 5.2.5 we will see that prescribing the slopes of the edges and the 1-chords of a polygon up to projective equivalence suffices for prescribing the whole polygon.

This construction only works for centrally symmetric polytopes (or polytopes projectively equivalent to them). In order to prescribe arbitrary polygons we deform the given polygon so that it is has one long edge and is otherwise very flat. We then glue a copy of this polytope to this long edge such that the result is centrally symmetric. We can prescribe this polytope and then use vertex forgetters to forget the double vertices.

5.2.1 Constructing a Special Shor Normal Form

In order to interface with Richter-Gebert's universality theorem (Theorem 4.13) we need to use Shor normal forms.

If a polynomial equation $f(x_1, \dots, x_n)$ is of the form $x_i + x_j - x_k = 0$ we say it is a *simple addition*. It is a *simple multiplication* if it is of the form $x_i \cdot x_j - x_k = 0$. If it is of one of the two forms we call it a *simple equation*. A *simple inequality* is one of the form $x_i - x_j < 0$ or $1 - x_i < 0$.

We remind the reader that a polynomial system in the variables $x_1, \dots, x_n$ is said to be in *Shor normal form* [46] if the only inequalities are simple inequalities establishing a total order of the variables which are all greater than one (i.e. $1 < x_1 < \dots < x_n < \infty$) and if the only equations are simple equations. Shor [54] proved that for every polynomial system there is a Shor normal form with an equivalent solution space. Equivalent here means *stably equivalent* [46]. Roughly speaking, the solution spaces of two polynomial systems are stably equivalent if they can be obtained by rational changes of variables and special, so-called *stable*, projections (eliminating variables with special restrictions). This is an algebraic notion which preserves the topology of the solution space, it implies homotopy equivalent.

We need a slight strengthening of his result: for every polynomial system such that the first m variables are totally ordered there is a Shor normal form also containing these m variables. Furthermore these two systems are equivalent in the way that the projections to these first m variables of the solution spaces of the two systems are equal.

Lemma 5.2 *Suppose we are given a polynomial system S (with algebraic coefficients) in the variables $x_1, \ldots, x_n$, with the additional restriction that for the first m variables we have $1 < x_1 < \ldots < x_m$.*

Then there is a system S' of polynomial equations of variables $x_1, \ldots, x_{n'}$ (with $n' \geq n$) in Shor normal form such that the following partial solution equivalence *holds:*

For each assignment for the first m variables $x_1 = \alpha_1, \ldots, x_m = \alpha_m$ we have the equivalence: the assignment can be completed to a solution of S if and only if it can also be completed to a solution of S'.

We will use this to transform a polynomial system in ordered variables $1 < x_1 < \ldots < x_m$ to a Shor normal form with auxiliary variables (all greater than x_m). This Shor normal form we can prescribe using the universality theorem. Edge forgetters will get rid of the auxiliary variables (see Section 5.2.2).

PROOF: Our proof of this lemma is a small extension to Shor's proof that every polynomial (in)equality system is equivalent to a system in Shor normal form. Our (and his) proof is divided into three steps: In each step we transform the polynomial system to another one which is more like the Shor normal form. Each of his transformation guarantees that the solutions of the systems can be translated one to the other. For each step we will give an extension of Shor's system which guarantees the partial solution equivalence.

In the first step Shor reduces the original system to a system containing only simple equations (simple additions and simple multiplications) and simple inequalities. He does this by introducing all intermediate steps of the polynomials as new variables.

In a slight modification of his process we first encode the coefficients of the polynomials in auxiliary variables which are defined by simple equations. We start with V_0 whose value is defined in the simple addition $V_0 = V_0 + V_0$ and V_1 which is defined by $V_1 = V_1 \cdot V_1$. We can then build powers of 2, $V_2 = V_1 + V_1$ and $V_{2^{i+1}} = V_{2^i} \cdot V_2$. We can now build all positive integers following the binary encoding: for a number n with the binary decomposition $n = \sum_{i=0}^{\lceil \log n \rceil} n_i 2^i$ this number is build according to this sum of products. Canonically we build negative integers $V_{-n} + V_n = V_0$ and rational numbers $V_{n/m} \cdot V_m = V_n$.

We also successively build the powers of variables: from x we get to V_{x^2}, an auxiliary variable representing the value x^2, by the equation $x \cdot x = V_{x^2}$. Again, we first encode the powers of a variables with an exponent which is a power of 2 and then by binary encoding all other powers of this variable.

By multiplying coefficients with powers of variables we get auxiliary variables for the monomials, which can be added to get (rationals) polynomials. Coefficients which are algebraic numbers can also be built now: Let α be an algebraic number which is defined by $f(\alpha) = 0$ for a rational polynomial and singled out from the other roots of f by the bounds $l < \alpha < r$ for rational numbers l and u. We have seen to define a variable V_α which is the solution for f and lies between l and r. This we can use as a coefficent now.

For example the system consisting of two variables x_1 and x_2 with $m = 1$, i.e. $x_1 > 1$, and the single inequalitiy $2x_1^2 - x_2 > \sqrt{2}$ is translated to the system consisting of seven variables V_1, V_2, $V_{\sqrt{2}}$, V_{x_1}, $V_{x_1^2}$, $V_{2x_1^2}$, V_{x_2}, and $V_{2x_1^2-x_2}$ and of the equations and inequalities

$$\begin{array}{rclrcl}
V_1 \cdot V_1 & = & V_1 & V_1 + V_1 & = & V_2 \\
V_{\sqrt{2}} \cdot V_{\sqrt{2}} & = & V_2 & V_1 & < & V_{\sqrt{2}} \\
V_{x_1} \cdot V_{x_1} & = & V_{x_1^2} & V_2 \cdot V_{x_1^2} & = & V_{2x_1^2} \\
V_{2x_1^2 - x_2} + V_{x_2} & = & V_{2x_1^2} & V_{\sqrt{2}} & < & V_{2x_1^2 - x_2} \\
V_1 & < & V_{x_1} & & &
\end{array}$$

Of course, after the introduction of auxiliary variables the variables $1 < x_1 < \ldots < x_m$ are still present in the (in)equality system and therefore the partial solution property holds.

In the second step Shor assumes a system S_{simple} in (relabeled) variables x_i consisting only of simple equations and inequalities. He constructs a new system where all variables are shifted by a large number a as to make them greater than 1: He introduces a variable V_a and variables V_{x_i+a} for each variable x_i (and more auxiliary variables) and constructs a system S_a of simple equations and inequalities which guarantees the following translation of solutions: For each solution of S_a the value of $x_i = V_{x_i+a} - V_a$ is a solution of the original system S_{simple}. Conversely, for each solution of S_{simple} and each number a the assignment $V_a = a$ and $V_{x_i+a} = x_i + a$ constitutes a solution of S_a.

As an example for this construction we note that the simple inequality $x_i < x_j$ becomes just $V_{x_i+a} < V_{x_j+a}$ and the simple addition $x_i + x_j = x_k$ becomes:

$$\begin{aligned} V_{x_i+a} + V_{x_j+a} &= V_{x_i+x_j+2a} \\ V_{x_k+a} + V_a &= V_{x_i+x_j+2a} \end{aligned}$$

The simple multiplication is more complicated, we refer to Shor's proof [54].

This translation works for all values of a and each variable is of the form $V_{\dots+ka}$ for some positive integer k, so Shor can require that all new (shifted) variables be greater than 1 and above translation still works for numbers a which are greater than the largest absolute value of a variable x_i. By the additional requirements that all $V's$ are greater than 1 we obtain the system $S_a^{>1}$.

We augment Shor's system $S_a^{>1}$ by introducing the variables V_{x_i} and the equations $V_{x_i} + V_a = V_{x_i+a}$ for all $i = 1, \dots, m$ (the first m variables). A solution of this system S_a^{new} translates to a solution of S_{simple} because it only added constraints to $S_a^{>1}$. On the other hand, in each solution of $S_a^{>1}$ the first m variables $V_{x_1+a}, \dots, V_{x_m+a}$ were such that $V_{x_i+a} - V_a$ was a solution for S_{simple}, hence greater than 1, so the requirement $V_{x_i} > 1$ is no new obstruction. Therefore the natural translation between solutions of $S_a^{>1}$ and S_a^{new} is valid. The systems S_{simple} and S_a^{new} are partial-solution equivalent since the ambiguity of the solutions of $S_a^{>1}$ introduced by adding the number a has cancelled out in S_a^{new} by subtracting it.

In the third step Shor translates a system of simple equations and inequalities in variables all greater than 1 (relabeled to x_i) to a system S_b: He shifts each variables by a different power of a large number b. He introduces for each variable x_i a variable $V_{x_i+b^i}$ (as well as a variable V_b and more auxiliary variables) and he gives simple equations and a *total order* of these variables such that the following equivalence holds: For all solutions of S_a^{new} there is a large number b_0 (for instance $b_0 = \max x_i + 1$) such that for all numbers $b > b_0$ the assignment $V_b = b$, $V_{x_i+b^i} = x_i + b^i$ and the induced values for the auxiliary variables constitute a solution of S_b. Conversely, $x_i = V_{x_i+b^i} - V_{b^i}$ computes a solution for S_a^{new} from a solution for S_b.

The only easy translation is from the simple inequality $x_i < x_j$: Shor introduces a new "unused" power of b, say b^α and writes

$$\begin{aligned} V_{x_i+b^i} + V_{b^\alpha - b^i} &= V_{x_i+b^\alpha} \\ V_{x_j+b^j} + V_{b^\alpha - b^j} &= V_{x_j+b^\alpha} \\ V_{x_i+b^\alpha} &< V_{x_j+b^\alpha} \\ V_{b^i} + V_{b^\alpha - b_i} &= V_{b^\alpha} \\ V_{b^j} + V_{b^\alpha - b_j} &= V_{b^\alpha} \end{aligned}$$

(Multiplication equations build the powers of b successively.) More auxiliary variables and equations are necessary for addition and multiplication. But since b is assumed to be very large (at least larger than any x_i), all these introduced variables can be brought into a total order, for example the above variables are ordered like:

$$V_{b^i} < V_{x_i+b^i} < V_{b^j} < V_{x_j+b^j} < V_{b^\alpha-b^j} < V_{b^\alpha-b^i} < V_{b^\alpha} < V_{x_i+b^\alpha} < V_{x_j+b^\alpha}.$$

The nice thing is that in the total order of the variables in S_b the variable V_b is the smallest one. So we can again augment this system by variables V_{x_i} for the first m variables $x_1, \ldots, x_m$ fitting in the total order *before* V_b:

$$1 < V_{x_1} < \ldots < V_{x_m} < V_b < \ldots$$

such that S_a^{new} and this system S_b^{new} are partial solution equivalent: Again the ambiguity of the shift by the b^i is canceled out by this extension, the proof of the validity of this operation is left to the reader. □

Corollary 5.3 *Let $\alpha_1 < \ldots < \alpha_m$ be algebraic numbers all greater than* 1. *Then there is a Shor normal form $S = S(\alpha_1, \ldots, \alpha_m)$ in variables $x_1, \ldots, x_n$ (with $n > m$) which has a solution and such that for all solutions of S the first m variables are $x_1 = \alpha_1, \ldots, x_m = \alpha_m$.*

PROOF: Let (f_i, l_i, u_i) the triples defining the α_i, i.e. f_i is a polynomial with integer coefficients such that α_i is the only root between the rational numbers l_i and u_i. The polynomial system $1 < x_1 < \ldots < x_m$, $f_i(x_i) = 0$, $l_i < x_i < u_i$ for all $i = 1, \ldots, m$ then has the unique solution $x_i = \alpha_i$. The Shor normal form constructed in Lemma 5.2 is partial solution equivalent with this polynomial system, hence it has a solution and each solution has the first m variables assigned to α_i. □

5.2.2 A Polytope Prescribing Algebraic Edge Slopes

In Section 4.3 we have seen that Richter-Gebert's universality theorem implied that for each Shor normal form S there is a polytope $P(S)$ such that every realization of $P(S)$ contains a normal computation frame encoding a solution of S and, conversely, every normal computation frame encoding a solution of S can be completed to a realization of $P(S)$ (Theorem 4.13).

We want to construct a combinatorial polytope $P(\alpha_1, \ldots, \alpha_m)$ which exactly prescribes the slopes of a normal 2-face to algebraic numbers $0 < 1 < \alpha_1 < \ldots < \alpha_m < \infty$:

By Lemma 5.3 we can translate the defining equations and inequalities of $1 < \alpha_1 < \ldots < \alpha_m$ into a Shor normal form $S(\alpha_1, \ldots, \alpha_m)$ in the variables $1 < x_1 < \ldots < x_n$ such that the α_i are the sole solutions of $x_1, \ldots, x_m$. Then $P(\alpha_1, \ldots, \alpha_m)$ can be constructed from the polytope constructed in Theorem 4.13, i.e. from $P(S(\alpha_1, \ldots, \alpha_m))$, by forgetting all edge pairs corresponding to the auxiliary variables $x_{m+1}, \ldots, x_n$ (see Figure 5.1).

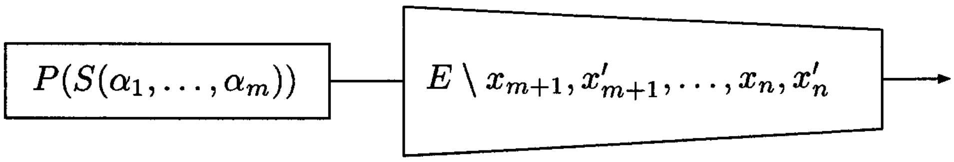

Figure 5.1: *Construction of* $P(\alpha_1, \ldots, \alpha_m)$

Lemma 5.4 *The polytope $P(\alpha_1, \ldots, \alpha_m)$ prescribes that one of its polygonal faces G prescribes that G is a normal computation frame and that G encodes the assignment $x_1 = \alpha_1, \ldots, x_m = \alpha_m$. The face G is the face of a pyramid facet.*

The last phrase of the lemma is important so that we can glue more stuff to (the pyramid over) this polygon and thereby impose more properties on G.

Algebraic numbers are fine, but we want to prescribe actual polygons. Let $\boldsymbol{G}$ be a (realized) normal polygon with algebraic vertex coordinates. Then the coordinates of the slopes $s(i) = i \wedge i'$ for $i = 1, \ldots, n$ of $\boldsymbol{G}$ are algebraic numbers. So are the numbers $\alpha_i = cr(s(1), s(2)|s(3), s(i+3))$ for $1 \le i \le m = n - 3$. These numbers are totally ordered and greater than 1. These are the numbers that are encoded in $\boldsymbol{G}$ if we view it as a computation frame where line 1 corresponds to ∞, line 2 corresponds to 0, line 3 corresponds to 1 and lines $i + 3$ correspond to the variables x_i. We define $P^{sl}(\boldsymbol{G}) = P(\alpha_1, \ldots, \alpha_{n-3})$. From the above lemma (and from Lemma 4.15 we conclude:

Corollary 5.5 *The combinatorial 4-polytope $P^{sl}(\boldsymbol{G})$ has a polygonal face G which is prescribed to be normal and have slopes which are projectively equivalent to the slopes of $\boldsymbol{G}$. The face G is the face of a pyramid facet.*

5.2.3 The Vertex Forgetter

We have seen how to presribe the slopes of edges of a polygonal face. In order to prescribe the exact shape of a polygon we will prescribe also the slopes of diagonals (chords) of the polygon. In order to translate prescribing of edge slopes to prescribing of chord slopes we want to "forget vertices." To this end we introduce the vertex forgetter polytope.

The *vertex forgetter* works quite similar to the edge forgetter: Start with a (realized) polygon $G(1,\dots,n)$, construct a prism over this polygon such that the edges connecting top and bottom faces meet in a point (their supporting lines, that is). But then consider the convex hull of the bottom polygon and all the vertices of the top polygon *except one*, say $i \wedge i+1$. Now perform a Lawrence construction analoguous to the edge forgetter, i.e. with q the the intersection point of the upright edges.

Let us list the facet sets as seen from q: $\mathcal{F}^0$ are the side faces of the prism, $\mathcal{F}^+$ the bottom polygon, $\mathcal{F}^-$ the top polygon and the new triangular facet. The upright edges constitute the set $\mathcal{R}^0$. Now the resulting polytope $V \setminus i \wedge i+1$ prescribes that the *vertices* of the top polygon are projectively equivalent to the corresponding vertices of the bottom polygon: The top polygon has *forgotten* the vertex $i \wedge i+1$. Note that since the facet sets can be seen purely combinatorially, the shape of the bottom and top polygons are not further prescribed.

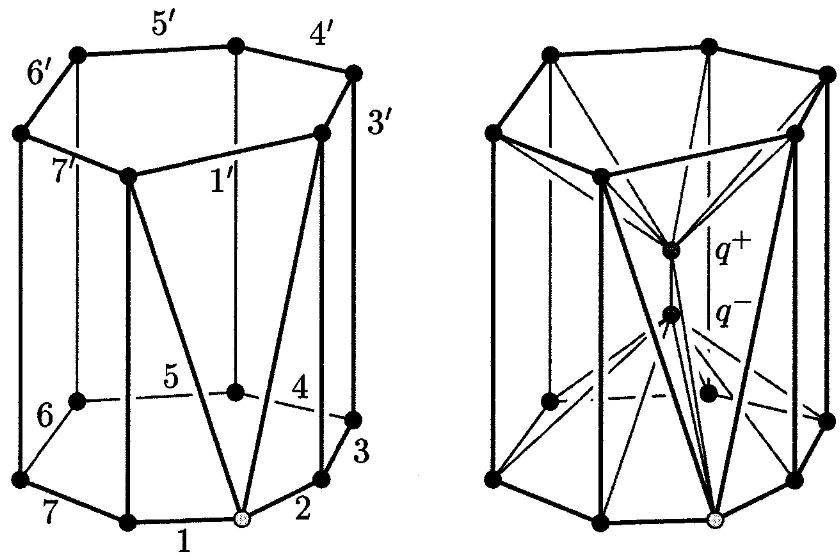

Figure 5.2: *Vertex forgetter* $V \setminus 1 \wedge 2$

As with the edge forgetters we can stack vertex forgetters on top of each other (by gluing along the pyramids) and forget many vertices $v_1, \ldots, v_r$. The polytope is then called $V \setminus v_1, \ldots, v_r$. Its symbol for the polytope diagrams is shown in Figure 5.3. The small end denotes the pyramid containing the polygon where the vertex has been forgotten.

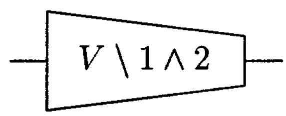

Figure 5.3: *Symbol for the vertex forgetter*

Sometimes it is better to regard this polytope as a *vertex inventor* this is especially the case when the small end is prescribed to have some shape, and other end has more vertices whose only restriction is that the result must be convex. The next section shows an application of this concept.

5.2.4 Prescribing Edge and 1-Chord Slopes of Centrally Symmetric Polytopes

Let $G^\diamondsuit$ be a centrally symmetric polytope with algebraic vertex coordinates. Let its edges be $1, 2, \ldots, n, 1', 2', \ldots, n'$ and suppose $n \geq 5$, i.e. the polytope has at least 10 edges. Figure 5.4 shows a projective transform of this polytope where the line l_∞ where all opposite edges and diagonals meet is finite.

We will construct a polytope which prescribes the edge slopes and the 1-chord slopes of one polygonal face to be projectively equivalent to the ones of $G^\diamondsuit$.

By applying Corollary 5.5 we can prescribe the slopes of the edges to be projectively equivalent to the slopes of $G^\diamondsuit$. By removing vertex $i \wedge i+1$ and its opposite vertex $i' \wedge i+1'$ we get a smaller polygon $G_i^\diamondsuit$ which is also normal and which has the (former) 1-chords $(i, i+1)$ and $(i', i+1')$ on its boundary. (Note that we define the index addition cyclically: $n+1 = 1'$ and $n+1' = 1$.)

Call $P^{sl}(G_i^\diamondsuit)$ the polytopes which prescribe the of a polygonal face to be the slopes of $G_i^\diamondsuit$ (again Corollary 5.5). Figure 5.5 shows how we glue these $P(G_i^\diamondsuit)$ together: On the right-hand side we see many connector polytopes which all share a $2n$-gon; eventually this will be prescribed to be projectively equivalent to $G^\diamondsuit$. The polytope $P^{sl}(G^\diamondsuit)$ on top assures that these $2n$-gons projectively have the same slopes as $G^\diamondsuit$. For each $i = 1, \ldots, n$ we glue a vertex-forgetter polytope to a connector, they forget the vertices $i \wedge i+1$ and

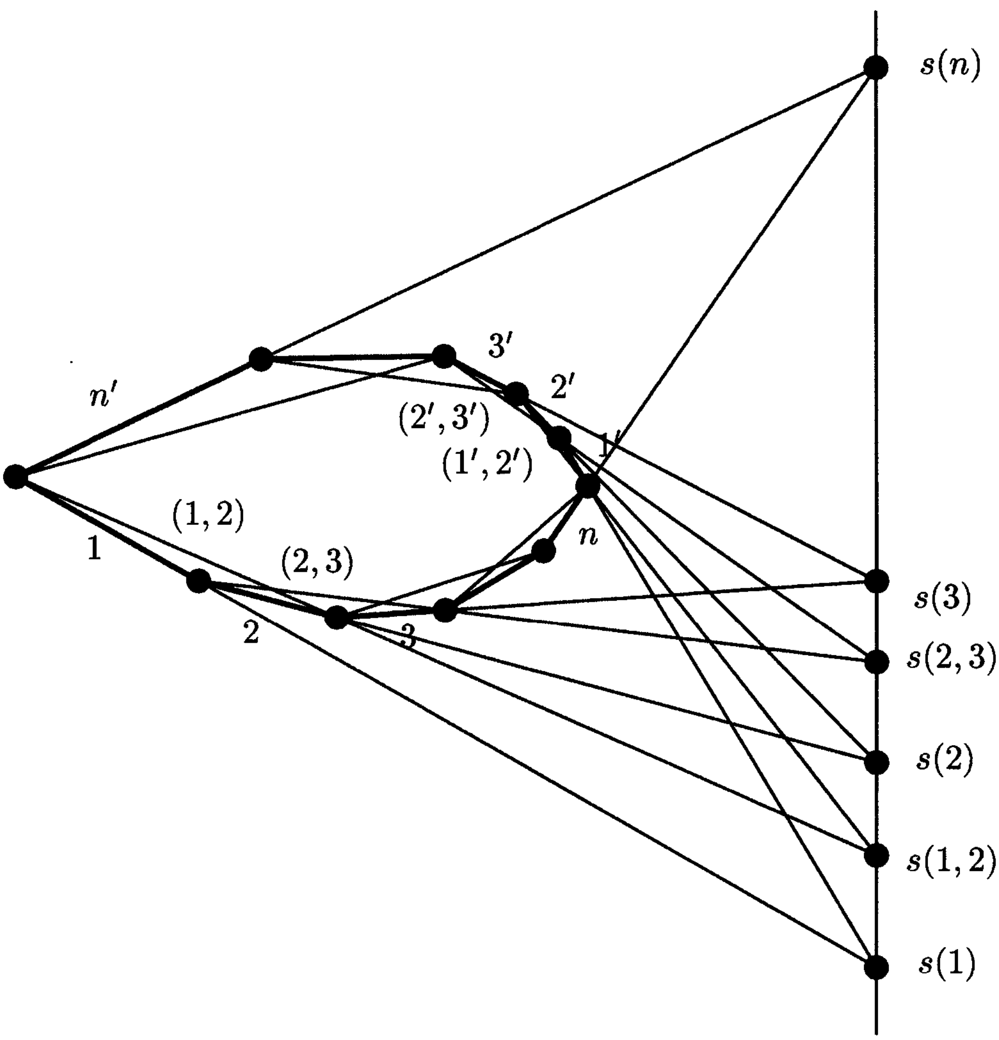

Figure 5.4: *Labeling in the centrally symmetric polygon* $G(1,2,\ldots,n,1',2',\ldots,n')$

$i' \wedge i+1'$. On the other side of this vertex forgetter we glue the $P^{sl}(G_i^\diamond)$. We call this polytope $P(G^\diamond)$.

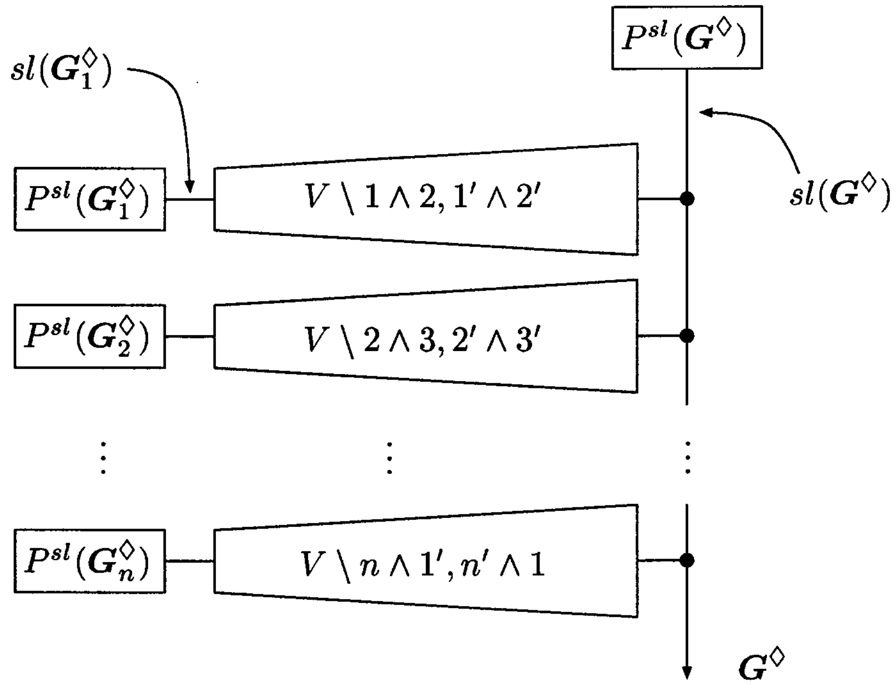

Figure 5.5: *Gluing diagram of the polytope* $P(G^\diamond)$

Lemma 5.6 *The polytope $P(G^\diamond)$ prescribes that the polygonal face in the last connector is normal with line l_∞, that opposite 1-chords also meet on the line l_∞, and that these points, the edge slopes and the 1-chord slopes, are projectively equivalent to the edge slopes and 1-chord slopes of $G^\diamond$. This polygonal face is the ground facet of a pyramid facet of $P(G^\diamond)$.*

PROOF: Consider a realization of the polytope $P(G^\diamond)$. The polygon in the connecting pyramid of $P^{sl}(G^\diamond)$ and the connector that is glued to it is normal and its edge slopes are projectively equivalent to the ones of $G^\diamond$. Denote it $sl(G^\diamond)$. Similarly denote $sl(G_i^\diamond)$ the polygons in the connecting pyramid of $P^{sl}(G_i^\diamond)$ (for $i = 1, \dots, n$) and the vertex forgetter which is glued to it. It is also normal and has edge slopes which are projectively equivalent to the ones of $G_i^\diamond$.

Consider $sl(G_1^\diamond)$. The vertex forgetters ensure that there is a projective transformation T_1 that map all vertices of $sl(G_1^\diamond)$ to the corresponding vertices of

$sl(G^{\diamond})$. Hence also in $sl(G^{\diamond})$ the the 1-chords $(1,2)$ and $(1',2')$ meet on l_{∞} (projective transformations commute with the meet operation)

Since only the vertices $1 \wedge 2$ and $1' \wedge 2'$ of $sl(G^{\diamond})$ are forgotten by the vertex forgetter the transformation T_1 maps the edges $3,4,5,\ldots,n$ and $3',4',5',\ldots,n'$ to their corresponding edges in $sl(G^{\diamond})$. Hence also $s(3)$, $s(4)$, and $s(5)$ and therefore all intersection points on the line l_{∞} are mapped to their correspondents: this follows from Lemma 4.14 since the slopes $s(3)$, $s(4)$, and $s(5)$ are a projective basis of l_{∞}. But this means that the projective equivalence of the lines l_{∞} of the polygons $G^{\diamond}$ and $sl(G^{\diamond})$ extends also to the slope of 1-chord $(1,2)$. The same argument works for all other 1-chord slopes $s(2,3),\ldots,s(n,1')$. This concludes the proof. □

5.2.5 Edge Slopes and 1-Chord Slopes Suffice

The next lemma implies that if all edge and 1-chord slopes of a normal polygon are prescribed then the polygon is prescribed (up to projective equivalence). The lemma makes a stronger statement: the polygon need not be normal, but the intersections of the edges and the 1-chords with an exterior line are prescribed. These intersection points we also call slopes (following the intuition of thinking of this line as the line at infinity).

Also it is not even necessary to know *all* 1-chord slopes: if all but two are known, the polygon is already prescribed. This strengthening of the statement will be very important in the next chapter when we reuse this lemma.

Lemma 5.7 *Let $G = G(1,\ldots,n)$ an n-gon with $n \geq 4$. Let l be a line outside of G. Furthermore let $l(i) = i \wedge l$ be the intersection of the line containing the edge i with l and let $l(i,j) = (i,j) \wedge l$ be the intersection containing the chord $(i,j) = (i-1 \wedge i) \vee (j \wedge j+1)$ with l. Let furthermore k_1 and k_2 two distinct edges (see Figure 5.6). Let $G = G(1',\ldots,n')$ be another n-gon, l' be a line outside G' and define $l'(i')$ and $l'(i',i+1')$ accordingly.*

If there is a projective transformation $f : l \to l'$ such that $f(l(i)) = l'(i')$ for all $i \in \{1,\ldots,n\}$ and $f(l(i,i+1)) = l'(i',i+1')$ for all $i \in \{1,\ldots,n\} \setminus \{k_1,k_2\}$ then G and G' are projectively equivalent.

The immediate consequence of this lemma and Lemma 5.6 is

Corollary 5.8 *Let $G^{\diamond}$ be a centrally symmetric polytope and $P(G^{\diamond})$ the combinatorial polytope which prescribes all edge and 1-chord slopes of a polygonal face G to be projectively equivalent to the edge and 1-chord slopes*

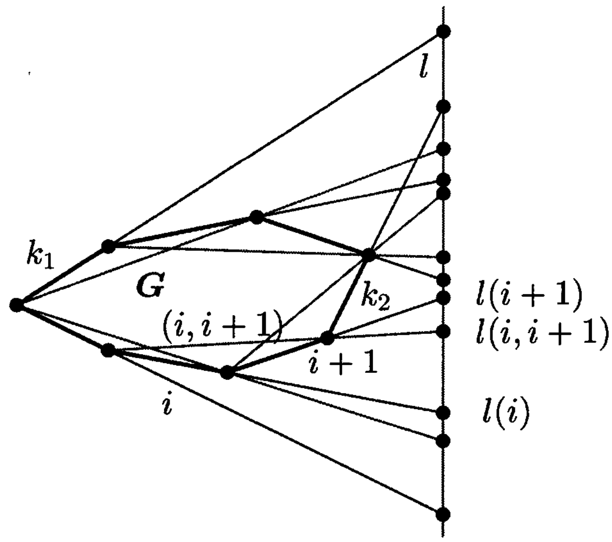

Figure 5.6: *Illustration of Lemma 5.7: if all intersection points with the exterior line l, i.e. the slopes, are known, G is determined up to projective equivalence*

of $G^{\diamond}$. The polytope $P(G^{\diamond})$ prescribes that G is projectively equivalent ot $G^{\diamond}$.

PROOF OF LEMMA 5.7: First we will show the following fact:

If the projective transformation f maps the slopes

$$\begin{array}{rcl} l(i) & \mapsto & l'(i') \\ l(i,i+1) & \mapsto & l'(i',i+1') \\ l(i+1) & \mapsto & l'(i+1') \\ \vdots & & \vdots \\ l(j-1,j) & \mapsto & l'(j-1',j') \end{array}$$

then it also maps $l(i,j) \mapsto l'(i',j')$.

We prove the fact using induction on $j-i$. The case $j-i=0,1$ are trivial. Suppose the fact is proven for chords up to difference $j-i<m$. We will now tackle the case $j-i=m$. Consider the following vertices: $i-1 \wedge i$, $i \wedge i+1$, $j-1 \wedge j$, and $j \wedge j+1$. They span the six lines (two edges and four chords) i, j, (i,j), $(i+1,j-1)$, $(i,j-1)$, $(i+1,j)$. The intersection points

of these lines with l are a quadrangular set (see Figure 5.7). Five of these are slopes map correctly under f since they are edge slopes or chord slopes with difference $< m$ (by induction assumption). So the sixth slope $l(i,j)$ must also be mapped to its counterpart $l'(i',j')$ (by Lemma 4.19 about quadrangular sets). This concludes the proof of the fact.

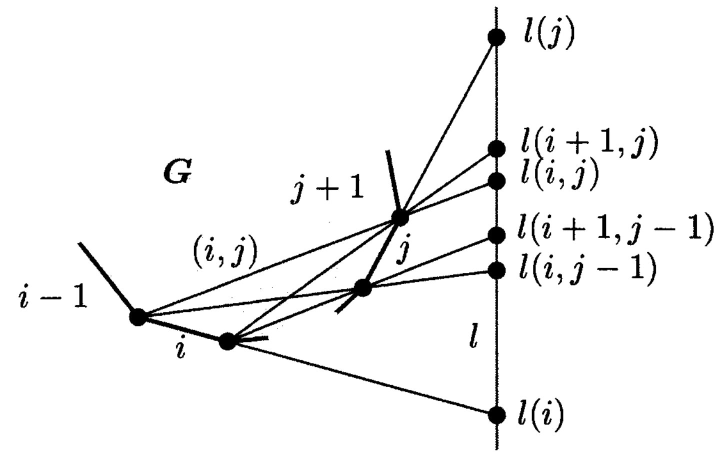

Figure 5.7: *The quadrangular set* $l(i)$, $l(j)$, $l(i,j)$, $l(i+1,j-1)$, $l(i,j-1)$, $l(i+1,j)$

The fact implies that $f(l(k_1+1,k_2)) = l'(k_1+1',k_2)$ since the edge slopes $l(i)$ for all i and the 1-chord slopes $l(i,i+1)$ for $k_1 < i \leq k_2$ are correctly mapped by f.

This is necessary in order to invoke Lemma 4.15: There is a unique projective transformation g which extends f and maps $k_1 - 1 \wedge k_1$ and $k_2 - 1 \wedge k_2$ to their counterparts in G'. We will show that $g(G) = G'$.

Notice that a vertex $i \wedge i+1$ between k_1 and k_2 (i.e. $k_1 < i < k_2$) is the intersection of the lines $l(k_1+1,i) \vee (k_1 \wedge k_1+1)$ and $l(i+1,k_2) \vee (k_2 \wedge k_2+1)$. Check that for the slopes $l(k_1+1,i)$ and $l(i+1,k_2)$ the assumptions of the fact at the beginning of the proof hold. Hence they are correctly mapped under $f = g|_l$. Since projective transformations commute with meet and join the vertex $i \wedge i+1$ is mapped correctly under g as well. The same kind of argument shows that the vertices $i \wedge i+1$ on the other side of k_1 and k_2 ($i < k_1$ or $k_2 < i$) are mapped correctly.

We have shown now that g maps all vertices of G correctly to their counterparts in G'. It remains to show that g is admissible for G, i.e. no point of G is mapped to infinity. Then we can invoke Lemma 4.21 to show that $g(G) = G'$. Since the preimage of infinity of g is a line, it suffices to show that no point on the boundary of G is mapped to infinity. No vertex is mapped to infinity, we have just shown that they are mapped to (finite) vertices.

Let us consider an edge i. The image of the line $\bar{i}$ contains the following four points: the vertices $i - 1' \wedge i'$ and $i' \wedge i + 1'$, the slope $l'(i')$ and the intersection with infinity $i \wedge \infty$. The last two points are not between the two vertices of the edge: l' was assumed to be outside of G' and of course an edge of a polygon does not contain infinite points. Hence by Lemma 4.20 the cross ratio $cr(i - 1' \wedge i', i' \wedge i + 1' | l'(i'), i' \wedge \infty) > 0$. Hence also for the preimages we have $cr(i - 1 \wedge i, i \wedge i + 1 | l(i), g^{-1}(i \wedge \infty)) > 0$ which again by Lemma 4.20 means that $g^{-1}(i \wedge \infty)$ lies not between $i - 1 \wedge i$ and $i \wedge i + 1$. Hence $g(i) = i'$ and therefore $g(G) = G'$. $\square$

5.2.6 Gluing Everything Together

Let G be a polygon with algebraic coordinates on $n \geq 6$ vertices. By a rational invertible projective transformation T we bring it into a *flat shape*, i.e. such that it has one long edge and other vertices are vertically above this edge. (This is not hard: consider two points a and b beyond an edge e of the polygon such that $a \vee b$ does not meet the polygon. The projective transformation mapping the endpoints of e back to themselves and a and b to different points at infinity not on $\bar{e}$ is such a transformation.)

Now we turn a second copy of the transformed polygon 180 degrees about the midpoint of this edge and take the union of these polygons (see Figure 5.8). We obtain a centrally symmetric polygon $G^\diamond$ on $2n - 2 \geq 10$ vertices. The coordinates of this polytope are obviously algebraic numbers.

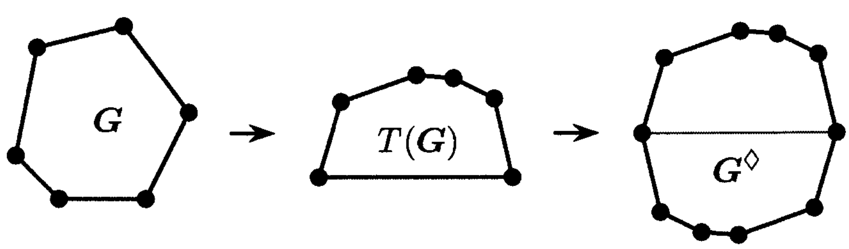

Figure 5.8: *Flattening and doubling a polygon gives a centrally symmetric polygon*

We can prescribe the centrally symmetric polytope $G^\diamond$. By deleting $n - 2$ vertices we can recover the flat polygon. We do this by gluing vertex forgetter

$V \setminus 1' \wedge 2', \ldots, n-1' \wedge n'$ to the prescribing polytope $P(G^{\diamond})$. Call the resulting polytope $P(G)$.

Corollary 5.9 *The polytope $P(G)$ prescribes its face to be projectively equivalent to G.*

5.3 $d+2$-Polytopes Prescribing d-Faces

This section is dedicated to showing that d-dimensional faces of $d+2$-polytopes can be exactly prescribed up to projective equivalence (for $d \geq 3$). Remember that the underlying idea about prescribing a polygonal face was to encode all information of the polygon in points on a line. Using Richter-Gebert's universality theorem we could "do arithmetic" with the points on this line.

This section follows this paradigm of encoding information about a polytope on lines:

Definition 5.10 *Let P be a polytope and e one of its edges. The* line image on e of a facet F in P *is the intersection of $\overline{e}$ with the hyperplanes $\overline{F}$ supporting F, i.e. $img_e(F) = e \wedge F$. The line image* of e in P *is the map*

$$\begin{array}{rcccl} img_e & : & facets(P) & \to & 2^{\mathbb{R}^d} \\ & & F & \mapsto & e \wedge F \end{array}$$

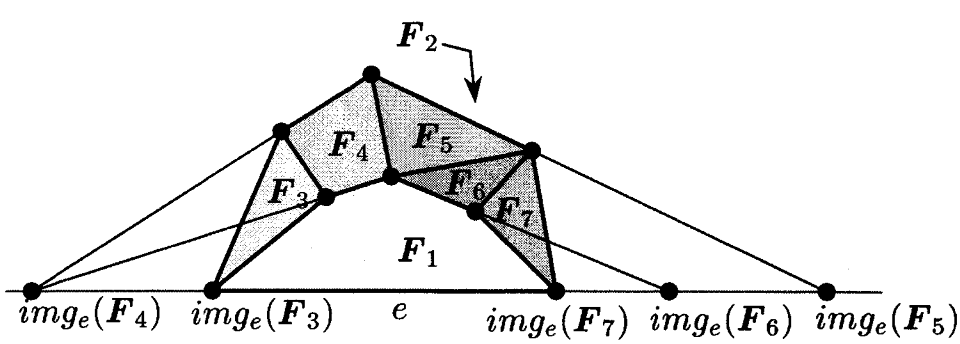

Figure 5.9: *The line image of an edge of the* taco polytope. *It is obtained by the convex hull of two polygons F_1 and F_2 sharing the edge e. Note that $\overline{e} = img_e(F_1) = img_e(F_2)$.*

Figure 5.9 shows a polytope $\boldsymbol{P}$ which is the convex hull of two polygons sharing an edge e and the line images on this edge e. Note that for the two facets sharing e the line image is the whole line. Note that line images can also be at infinity and that different facets can have the same line image.

The main fact leading to our result is: A d-dimensional polytope is determined up to projective equivalence if the line images of its are determined up to projective equivalence. More specifically, if all line images of edges in a certain edge set E of two combinatorially equivalent polytopes $\boldsymbol{P}$ and $\boldsymbol{P}'$ are projectively equivalent by possibly *different* projective transformations then there is *one* projective transformation taking $\boldsymbol{P}$ to $\boldsymbol{P}'$.

The construction of the combinatorial $d+2$-polytope which encodes the exact shape $\boldsymbol{P}$ of one of its d-faces is divided into three parts:

1. In Section 5.3.1 we will show Lemma 5.11 which encapsulates the main fact.

2. In Section 5.3.2 we will construct a $d+1$-polytope which has one facet combinatorially equivalent to $\boldsymbol{P}$ and a polygonal face G_e which encodes the line image of $\boldsymbol{P}$ of one edge e. Then we will show how to glue a $d+1$-polytope to this polytope which prescribes the exact shape of G_e, and therefore the line image with respect to the one edge.

3. Finally in Section 5.3.3, we will show how the these construction pieces act together. We do this construction for every edge of $\boldsymbol{P}$. We lift the combinatorial $d+1$-polytopes which prescribe one line imgage of $\boldsymbol{P}$ to $d+2$-space by erecting pyramids over it. We construct a $d+2$-dimensional connector polytope and use it to glue the one-edge-image prescribing polytopes together. Finally, we show how to perform slight changes to $\boldsymbol{P}$ in order to invoke the main fact. The resulting polytope will prescribe one of its d-faces to be projectively equivalent to $\boldsymbol{P}$.

Most of our constructions we will only show how to do for $d = 3$. This is quite feasible since the arguments easily extend to higher dimensions. In the cases where this extension is not canonical we will specifically note this.

5.3.1 Few Line Images Suffice

We will show now that under certain conditions on the face lattice of $\boldsymbol{P}$ (a vertex of degree d which is only called in simplex facets and which is "far away" from at least one facet) the line images of a few edges suffice to determine exact shape of the polytope.

Lemma 5.11 *Let $\boldsymbol{P}$ and $\boldsymbol{P}'$ be d-dimensional polytopes with identical face lattice. Furthermore, let p be a vertex of degree d of $\boldsymbol{P}$ such that*

1. *the neighbors of p are $q_0, \ldots, q_d$,*

2. *all facets incident to p are simplices, i.e. facets with vertex sets*

$$\{\overbrace{p, q_1, \ldots, q_d}^{\text{omit } p_i}\},$$

3. *there is a facet $\boldsymbol{G}$ of $\boldsymbol{P}$ such that $\boldsymbol{G} \cap \{p, q_1, \ldots, q_d\} = \emptyset$,*

4. *there are projective transformations $f_i : p \vee q_i \to p' \vee q_i'$ such that $f_i(img_{p \vee q_i}(\boldsymbol{F})) = img_{p' \vee q_i'}(\boldsymbol{F}')$ for all facets $\boldsymbol{F}$ of $\boldsymbol{P}$.*

Then there is a projective transformation f such that $f(\boldsymbol{P}) = \boldsymbol{P}'$. In particular, $f|_{p \vee q_i} = f_i$.

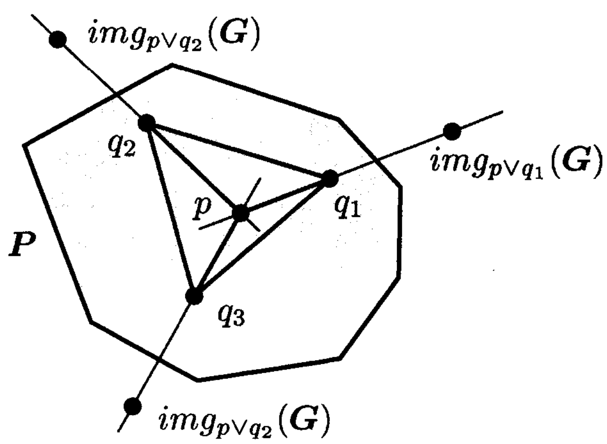

Figure 5.10: *Top view of $\boldsymbol{P}$, with the line images of $\boldsymbol{G}$ on the edges $p \vee q_i$ ($i = 1, 2, 3$)*

PROOF: Whenever we consider some face of $\boldsymbol{P}$, by attaching a prime to its name we mean image under the face lattice isomorphism to $\boldsymbol{P}'$. For instance, if $\boldsymbol{F}$ is a face of $\boldsymbol{P}$, then $\boldsymbol{F}'$ is the corresponding face of $\boldsymbol{P}'$.

First, we will now construct a candidate projective transformation f, then that it really has the desired properties: Since $\boldsymbol{G}$ does not contain any of

$q, q_1, \ldots, q_d$, by convexity $\overline{G}$ does not either. Furthermore by the face structure of $\boldsymbol{P}$ the $p, q_1, \ldots, q_d$ are in general position. Hence by Lemma 4.17 there is a unique projective transformation f such that

$$\begin{aligned} f(\overline{G}) &= \overline{G'}, \\ f(p) &= p', \\ f(q_i) &= q_i' \qquad \text{for all } i = 1, \ldots, d. \end{aligned}$$

Since p, q_i, and $(p \vee q_i) \wedge \boldsymbol{G}$ are distinct points on $p \vee q_i$ and since f_i and f have the same value at these points, by Lemma 4.14 they must be equal:

$$f|_{p \vee q_i} = f_i.$$

We will now show that f maps all supporting hyperplanes of facets to the supporting hyperplanes of the corresponding facets. It follows that the vertices are mapped of $\boldsymbol{P}$ are mapped to the corresponding vertices of $\boldsymbol{P}'$. Eventually, we will show that f is admissible, i.e. it maps $\boldsymbol{P}$ to $\boldsymbol{P}'$.

Consider the facets $\boldsymbol{F}_i = \overbrace{p \vee q_1 \vee \ldots \vee q_d}^{\text{omit } q_i}$. The transformation f maps their supporting hyperplanes to the corresponding supporting hyperplane of $\boldsymbol{P}'$.

Let $\boldsymbol{F}$ be one of the remaining facets of $\boldsymbol{P}$. None of the points $(p \vee q_i) \wedge \boldsymbol{F}$ coincides with p since $\boldsymbol{F}$ is not incident to p. From $p, q_1, \ldots, q_d$ being in general position, it follows that also $(p \vee q_i) \wedge \boldsymbol{F}$ are in general position, hence span $\overline{\boldsymbol{F}}$. The transformation f takes the line image of $p \vee q_i$ to the line image of $p' \vee q_i'$ (on $p \vee q_i$ it coincides with f_i). So f maps a set of points spanning $\overline{\boldsymbol{F}}$ to points in $\overline{\boldsymbol{F}'}$. Since f is invertible, $f(\overline{\boldsymbol{F}}) = \overline{\boldsymbol{F}'}$.

It remains to show that f is also admissible, i.e. that it maps the convex hull of the vertices of $\boldsymbol{P}$ to the convex hull of the vertices of $\boldsymbol{P}'$. If f was not admissible, by Lemma 4.21 it would map points inside $\boldsymbol{P}$ to infinity. If $f^{-1}(\infty)$, the preimage of the hyperplane at infinity, intersected $\boldsymbol{P}$, then there would also be an edge intersecting this hyperplane: In this case, since $f^{-1}(\infty)$ partitions the vertices of $\boldsymbol{P}$ according to its sides and since the vertex-edge graph of $\boldsymbol{P}$ is connected, there will be an edge whose endpoints are on both sides of $f^{-1}(\infty)$. Hence it remains to show that no edge $r \vee s$ of $\boldsymbol{P}$ intersects $f^{-1}(\infty)$.

For each edge $r \vee s$ of $\boldsymbol{P}$ there is a facet $\boldsymbol{F}$ which contains neither r nor s (i.e. r, s, and $(r \vee s) \wedge \boldsymbol{F}$ are three distinct points): For the edges having both endpoints in $\{p, q_1, \ldots, q_d\}$ the facet $\boldsymbol{G}$ takes this role. For all other edges, at most one of r and s is equal to some q_i. Then $\boldsymbol{F}_i$ is a facet not containing r and s.

The facet $\boldsymbol{F}$ is such that point $(r \vee s) \wedge \boldsymbol{F}$ is outside of the edge $r \vee s$. The same holds in $\boldsymbol{P}'$, i.e. $(r' \vee s') \wedge \boldsymbol{F}'$ lies outside of $r' \vee s'$. Of course the hyperplane at infinity meets $r' \vee s'$ in a point outside of the edge. Hence by Lemma 4.20 the cross ratio

$$cr(r', s'|(r' \vee s') \wedge \boldsymbol{F}', (r' \vee s') \wedge \infty) > 0.$$

This cross ratio is invariant under the projective transformation $f^{-1}|_{r' \wedge s'}$, hence

$$cr(r, s|(r \vee s) \wedge \boldsymbol{F}, (r \vee s) \wedge f^{-1}(\infty)) > 0.$$

It follows by Lemma 4.20 that $(r \vee s) \wedge f^{-1}(\infty)$ lies outside of the edge $r \vee s$. □

5.3.2 A $d+1$-Polytope Prescribing one Line Image

In order to construct a polytope which prescribes the line image of one of its facets we will proceed in two steps:

1. We will construct a $d+1$-polytope which links the line image an edge of one of its facets P to the line image of a polygon G (Section 5.3.2).

2. We will glue another polytope to this polytope which prescribes the shape of G, thereby prescribing the one line image of P (Section 5.3.2).

Polygon Slopes Linked to Line Image

Let P be the combinatorial polytope corresponding to the polytope $\boldsymbol{P}$ we want to prescribe and e an edge of P. We will now construct a combinatorial $d+1$-polytope P_e which has P as a d-face and a polygon G_e as a 2-face such that

1. the faces P and G_e share the edge e, and

2. in every realization $\boldsymbol{P}_e$ of P_e the line image of the d-face $\boldsymbol{P}'$ realizing P is encoded in intersections of the edges of the realization $\boldsymbol{G}_e$ of G_e. More precisely, for each facet $\boldsymbol{F}$ of $\boldsymbol{P}'$ there will be an edge i of $\boldsymbol{G}_e$ such that
$$img_e(\boldsymbol{F}) = e \wedge i.$$

Of course, the points $e \wedge i$ are the line image of $\boldsymbol{G}_e$, so we could say that we link the line image of P to the line image of G.

We will construct P_e by constructing a realized polytope $\boldsymbol{P}_e$ and defining P_e to be the combinatorial polytope corresponding to $\boldsymbol{P}_e$.

Let $\boldsymbol{F}_1$ and $\boldsymbol{F}_2$ be two facets of $\boldsymbol{P}$ incident to the edge e. Let the facets of $\boldsymbol{P}$ be partitioned into sets $\mathcal{F}_0, \ldots, \mathcal{F}_n$ such that two facets $\boldsymbol{F}$ and $\boldsymbol{F}'$ are in the same set if and only if their line images on e coincide, i.e. $img_e(\boldsymbol{F}) = img_e(\boldsymbol{F}')$. W.l.o.g. $\mathcal{F}_0$ consists of $\boldsymbol{F}_1$ and $\boldsymbol{F}_2$. Order the set $\mathcal{F}_1, \ldots, \mathcal{F}_n$ according to their occurence on the line image of e. For all $i = 1, \ldots, n$ denote $p_i = img_e(\boldsymbol{F})$ for $\boldsymbol{F} \in \mathcal{F}_i$. Hence p_1 and p_n are the endpoints of e. E.g. for the polytope in Figure 5.9 we have $p_1 = img_e(\boldsymbol{F}_7), p_2 = img_e(\boldsymbol{F}_6), \ldots, p_5 = img_e(\boldsymbol{F}_3)$.

Now embed $\boldsymbol{P}$ into $\mathbb{R}^{d+1}$ and consider a polygon $\boldsymbol{G}_e$ which lies in a plane that shares only the line supporting the edge e with the d-dimensional hyperplane $\overline{\boldsymbol{P}}$ and has the following properties: Besides e it has n edges $1, \ldots, n$ and for all edges $i = 1, \ldots, n$ the intersection of the supporting lines of i and e coincide with the line image p_i, i.e. $i \wedge e = p_i$.

Such a polygon is easy to construct: First draw edge 1 into a space direction that misses $\boldsymbol{P}$ (i.e. the projective 2-flat $e \vee 1$ intersects the hyperplane containing $\boldsymbol{P}$ only in $\overline{e}$). Then connect the endpoint of edge 1 with p_2 and construct edge 2 on this line such that the convex hull of e, 1, and 2 has theses three edges on their boundary: This is easy to achieve by letting edge 1 go off from the intersection with edge 0 in the right direction and by letting it be short enough. Construct the edges $3, \ldots, n-1$ in the same way, always taking care that (after having constructed edge i) the convex hull of $e, 0, \ldots, i$ has these edges on its boundary. At the end the edge n is already constructed. See Figure 5.11.

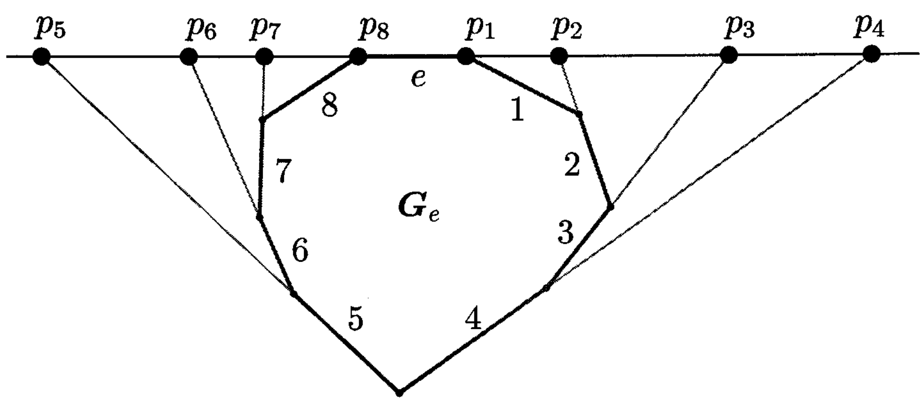

Figure 5.11: *The polygon $\boldsymbol{G}_e$. Attention: $\boldsymbol{G}_e$ does* not *lie in the same 3-dimensional hyperplane as $\boldsymbol{P}$*

Define $\boldsymbol{P}_e$ to be the convex hull of $\boldsymbol{P}$ and $\boldsymbol{G}_e$. The nice thing is now that the facets of this $d+1$-dimensional polytope can be combinatorially described if the facet partition $\mathcal{F}_i$ is known and that we can use the face lattice of $\boldsymbol{P}_e$ to prescribe the connection between the line image img_e of $\boldsymbol{P}$ and the polygon $\boldsymbol{G}_e$ already on the combinatorial level.

Lemma 5.12 *The polytope $\boldsymbol{P}_e = conv(\boldsymbol{P}, \boldsymbol{G}_e)$ has the following properties:*

1. *all vertices of $\boldsymbol{P}$ and of $\boldsymbol{G}_e$ are vertices of $\boldsymbol{P}_e$,*
2. *the polytope $\boldsymbol{P}$ is a facet of $\boldsymbol{P}_e$ and the polygon $\boldsymbol{G}_e$ is a face of $\boldsymbol{P}_e$,*
3. *the polytopes $conv(\boldsymbol{F}_1 \cup \boldsymbol{G}_e)$ and $conv(\boldsymbol{F}_2 \cup \boldsymbol{G}_e)$ are facets and they are necessarily flat,*
4. *for all $i = 1, \ldots, n$ and all $\boldsymbol{F} \in \mathcal{F}_i$ the polytopes $conv(\boldsymbol{F} \cup i)$ are facets.*

Remark: We know much more about the polytope $\boldsymbol{P}_e$. For instance, we have not listed all facets of $\boldsymbol{P}_e$: we omitted the tetrahedral facets conv$(R \cup j)$ for all edges R of $\boldsymbol{P}$ that are incident to two facets $\boldsymbol{F} \in \mathcal{F}_i$ and $\boldsymbol{F}' \in \mathcal{F}_k$ and $i < j < k$. (This completes the list of facets of $\boldsymbol{P}_e$.) This shows that the combinatorial structure of $\boldsymbol{P}_e$ is already known once the partition of the facets of $\boldsymbol{P}$ into the $\mathcal{F}_i$ is known.

Another Remark: Furthermore, in dimension $d = 3$ we know a (relatively) simple description of what these facets look like: The facets of the form conv$(\boldsymbol{F} \cup i)$ for facets $\boldsymbol{F} \in \mathcal{F}_i$ are Lawrence extensions of $\boldsymbol{F}$ with respect to the point p_i. The facet conv$(\boldsymbol{F}_1 \cup \boldsymbol{G}_e)$ has the following facets: (1) $\boldsymbol{F}_1$ and $\boldsymbol{G}_e$ themselves, (2) quadrilateral or triangular facets conv$(f \cup i)$ for edges of $\boldsymbol{F}_1$ such that $f = \boldsymbol{F}_1 \wedge \boldsymbol{F}$ for some facet $\boldsymbol{F} \in \mathcal{F}_i$ and (3) triangular facets conv$(v \cup j)$ for vertices v of $\boldsymbol{F}_1$ which are between two edges $f = \boldsymbol{F}_1 \wedge \boldsymbol{F}$ and $f' = \boldsymbol{F}_1 \wedge \boldsymbol{F}'$ with $\boldsymbol{F} \in \mathcal{F}_i$, $\boldsymbol{F}' \in \mathcal{F}_k$, and $i < j < k$. The facet $\boldsymbol{F}_2$ has an analoguous description. However, we do not make explicit use of these properties in the sequel, so we omit the easy but somewhat technical proof. In Figure 5.12 the reader can find all these properties at work.

PROOF OF LEMMA 5.12: The intersection of the d-dimensional hyperplane $\overline{P}$ and the 2-dimensional plane $\overline{\boldsymbol{G}_e}$ is the line $\overline{e}$. Therefore $\overline{P}$ contains none of the vertices of $\boldsymbol{G}_e$ except the vertices on e. All of the vertices of $\boldsymbol{G}_e$ lie on the same side of this hyperplane. Hence $\overline{P}$ supports the facet $\boldsymbol{P}$ of $\boldsymbol{P}_e$ and all of $\boldsymbol{P}$'s vertices are vertices of $\boldsymbol{P}_e$.

We want to prove now that vertices v of $\boldsymbol{G}_e$ are vertices of $\boldsymbol{P}_e$ (w.l.o.g. we can assume $v \notin e$. Let e_v be a supporting line of v and $p_v = e_v \wedge e$. Then p_v is

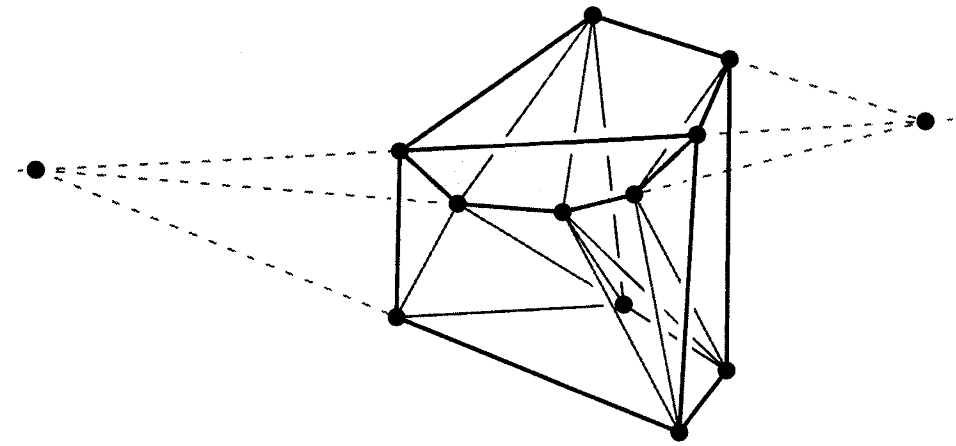

Figure 5.12: *Schlegel diagram of $\boldsymbol{P}_e$ where $\boldsymbol{P}$ is a (warped) cube: $\boldsymbol{F}_1$ = front facet, $\boldsymbol{F}_2$ = top facet, $\mathcal{F}_0 = \{$right facet$\}$, $\mathcal{F}_1 = \{$back facet$\}$, $\mathcal{F}_2 = \{$bottom facet$\}$, $\mathcal{F}_3 = \{$left facet$\}$*

outside of e, but on $\overline{e}$, therefore outside of $\boldsymbol{P}$. Hence there is a 2-dimensional hyperplane $H \subset \overline{\boldsymbol{P}}$ which contains p_v, but which does not cut $\boldsymbol{P}$. Then the 3-dimensional hyperplane $H \vee v$ contains v, but has all other points of $\boldsymbol{P}$ and $\boldsymbol{G}_e$ on the same side, i.e. it supports v.

The facet $\boldsymbol{F}_1$ of $\boldsymbol{P}$ is defined by a $d-1$-dimensional supporting plane. This plane contains e, so the span H of this plane with some vertex of $\boldsymbol{G}_e$ contains all of $\boldsymbol{G}_e$. The d-dimensional hyperplane H has all other vertices of $\boldsymbol{P}$ on its other side, hence H is a supporting hyperplane of $\boldsymbol{P}_e$ and defines the facet conv$(\boldsymbol{F}_1 \cup \boldsymbol{G}_e)$. By the same reasoning conv$(\boldsymbol{F}_2 \cup \boldsymbol{G}_e)$ is a facet of $\boldsymbol{P}_e$. The intersection of the two facets conv$(\boldsymbol{F}_1 \cup \boldsymbol{G}_e) \cap$ conv$(\boldsymbol{F}_2 \cup \boldsymbol{G}_e)$ is a face of $\boldsymbol{P}_e$. It contains all vertices of $\boldsymbol{G}_e$, but no other vertices of $\boldsymbol{P}$ than the endpoints of e. Therefore this face is $\boldsymbol{G}_e$ and all vertices of $\boldsymbol{G}_e$ are vertices of $\boldsymbol{P}_e$.

Let $\boldsymbol{F} \in \mathcal{F}_i$ be a facet of $\boldsymbol{P}$. Suppose $i < n$. Define the d-dimensional plane $H^i_{\boldsymbol{F}} = \boldsymbol{F} \vee (i \wedge i+1)$. We claim that $H^i_{\boldsymbol{F}}$ is a supporting hyperplane of $\boldsymbol{P}_e$. This hyperplane is d-dimensional since the hyperplane $\overline{\boldsymbol{F}} \subset \overline{\boldsymbol{P}}$ and $i \wedge i+1$ is not in this hyperplane, so $\overline{\boldsymbol{F}}$ cannot contain this endpoint. The hyperplane $H^i_{\boldsymbol{F}}$ contains p_i since $\overline{\boldsymbol{F}}$ does. Hence it also contains the other endpoint $i-1 \wedge i$ of i. Therefore it has all other vertices of $\boldsymbol{G}_e$ on one side. Since it contains $\overline{\boldsymbol{F}}$ it has all vertices of $\boldsymbol{P}$ on one side as well. These have to be the same sides since e is shared by $\boldsymbol{G}_e$ and $\boldsymbol{P}$. Hence conv$(\boldsymbol{F} \cup i)$ is a face. It is also a facet since $\boldsymbol{F}$ is $d-1$-dimensional and i is not contained in the hyperplane containing $\boldsymbol{F}$. The case $i = n$ works analoguously, when we exchange the roles of $i-1 \wedge i = n-1 \wedge n$ and $i+1 \wedge i = e \wedge n$. □

Let P_e be the combinatorial polytope with the same face lattice as $\boldsymbol{P}_e$, let P be its facet corresponding to the polytope $\boldsymbol{P}$ and G_e be the 2-dimensional face corresponding to $\boldsymbol{G}_e$. The previous lemma implies that in any realization of P_e the line image in the realization of P of the edge e is encoded in the realization of G_e.

Corollary 5.13 *Let $\boldsymbol{P}'_e$ be a realization of P_e, $\boldsymbol{P}'$ and $\boldsymbol{G}'_e$ the induced realizations of P and G_e. Then for $\boldsymbol{P}'$ the line image of the edge e is*

$$img_e(\boldsymbol{F}) = i \wedge e$$

for all facets $\boldsymbol{F} \in \mathcal{F}_i$ of $\boldsymbol{P}'$ and all edges $i = 1, \ldots, n$ of $\boldsymbol{G}'_e$.

PROOF: For a facet $\boldsymbol{F} \in \mathcal{F}_i$ the facet $\boldsymbol{F} \vee i$ spans a d-dimensional supporting hyperplane of $\boldsymbol{P}'_e$. This hyperplane meets $\bar{e}$ in exactly one point. But also $\overline{\boldsymbol{F}}$ and $\bar{i}$ meet $\bar{i}$ in exactly one point. It follows that all these points have to be equal, hence $\boldsymbol{F} \wedge e = i \wedge e$. □

Prescribing One Line Image

We will show now how to construct a combinatorial $d+1$-polytope which prescribes the line image of one edge of one of its facets.

First we attach pyramids to P_e in order to obtain a facet $\boldsymbol{pyr}(\ldots(\boldsymbol{pyr}(G_e))\ldots)$. Along this facet we will then glue pyramids over pyramids etc. over a 4-polytope prescribing the exact shape of G_e.

We erect a pyramid on P_e over its facet $F_1 \cup G_e$: By an abuse of notation we write $\boldsymbol{pyr}(F_1 \cup G_e)$ for the pyramid whose ground facet combintorially equivalent to the facet $F_1 \cup G_e$ of P_e. Consider the connected sum $P_e \#_{F_1 \cup G_e} \boldsymbol{pyr}(F_1 \cup G_e)$. The facet $F_1 \cup G_e$ is necessarily flat since in all realizations of the $d-1$-faces all vertices lie in the two hyperplanes spanned by the vertices in F_1 and in G_e, and these hyperplanes share the line $\bar{e}$ (Lemma 4.8). Since one face of $\boldsymbol{pyr}(F_1 \cup G_e)$ is G_e the new polytope has a face $\boldsymbol{pyr}(G_e)$.

In dimension $d = 3$ we are done since this face $\boldsymbol{pyr}(G_e)$ is already a facet. In higher dimensions, we erect a pyramid over a facet of this polytope containing the face $\boldsymbol{pyr}(G_e)$. The resulting polytope has a face $\boldsymbol{pyr}(\boldsymbol{pyr}(G_e))$. We proceed attaching pyramids until we end up with a facet$\boldsymbol{pyr}(\ldots(\boldsymbol{pyr}(G_e))\ldots)$. Note in all the later steps we glued along pyramids which are necessarily flat.

Along $\boldsymbol{pyr}(\ldots(\boldsymbol{pyr}(G_e))\ldots)$ we glue an iterated pyramid over a prescribing polytope $P(\boldsymbol{G}_e)$ of Section 5.2.6 which prescribes that the polygon is projectively equivalent to $\boldsymbol{G}_e$. (Note that $P(\boldsymbol{G}_e)$ is a 4-polytope which already has a facet $\boldsymbol{pyr}(G_e)$.) We obtain the combinatorial polytope $P_e^{pre}(\boldsymbol{P})$.

Since we glue only along necessarily flat faces and since iterated pyramids do not change properties prescribed by the polytope $P(\boldsymbol{G}_e)$ we obtain the following lemma:

Lemma 5.14 *Every realization $\boldsymbol{P}_e^{pre\prime}$ of $P_e^{pre}(\boldsymbol{P})$ has the property that its facet $\boldsymbol{P}'$ (which is combinatorially equivalent to $\boldsymbol{P}$) has a the line image on e which is projectively equivalent to the line image on e in $\boldsymbol{P}$.*

Note that we have not shown (but we could) that $P_e^{pre}(\boldsymbol{P})$ *prescribes* this property since not every realization of P with the right line image on e might extend to a realization of P_e. (After the statement of Lemma 5.12 we remarked that this is so, but we have not proven it.) However, our statement is sufficient for the proof of the main result.

5.3.3 Gluing many Line-Image Prescriptors

We have seen how to construct a $d+1$-dimensional polytope $P_e^{pre}(\boldsymbol{P})$ which has a facet P which is combinatorially equivalent to $\boldsymbol{P}$ such that in realization the line image w.r.t. the edge e is projectively equivalent ot the line image in $\boldsymbol{P}$ (Lemma 5.14). Also we have seen that under some conditions the line images of a few edges suffice to determine the exact shape of a polytope (Lemma 5.11).

We will now see how to glue together the polytopes $P_e^{pre}(\boldsymbol{P})$ for some edges e. First we will show how to augment a polytope $\boldsymbol{P}$ to a polytope $\boldsymbol{P}_{aug}$ which satisfies the assumptions of Lemma 5.11. Then we will show how to glue together the different prescriptor polytopes $P_e^{pre}(\boldsymbol{P}_{aug})$. It will become necessary to lift the construction to dimension $d+2$. At the end we will show how to rectify the augmentation such that the resulting polytope prescribes that a face is projectively equivalent to $\boldsymbol{P}$.

Augmenting the polytope $\boldsymbol{P}$

In order to be sure that the line images of some edges determine the exact shape of the polytope in Lemma 5.11 we needed the conditions that there is a degree-d vertex p which is only incident to facets which are simplices and such that there is a “far-away” facet $\boldsymbol{G}$.

Let $\boldsymbol{P}$ be the given d-polytope and $\boldsymbol{F}$ one of its facets. Consider d points $q_1, \dots, q_d$ beyond $\boldsymbol{F}$ of $\boldsymbol{P}$ which are in a hyperplane which is parallel to $\boldsymbol{F}$. It is not hard to see that those points are easy to construct (even in polynomial time). The convex hull conv($\boldsymbol{P} \cup \{q_1, \dots, q_d\}$) thus has a simplex facet $q_1 \vee \dots \vee q_d$. Let p be a vertex beyond this facet. We define $\boldsymbol{P}_{aug}$ be the convex hull conv($\boldsymbol{P} \cup \{p, q_1, \dots, q_d\}$) (see Figure 5.13).

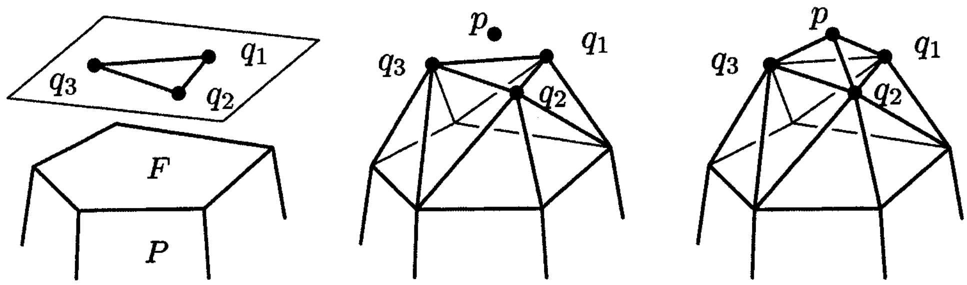

Figure 5.13: *Construction of p and* $q_1, \dots, q_d$

Obviously, p is only adjacent to the vertices q_i. Also the facets that contain p are exactly $p \vee q_1 \vee \dots \vee q_{i-1} \vee q_{i+1} \vee \dots q_d$. All of $\boldsymbol{P}$'s facets except $\boldsymbol{F}$ are still facets of $\boldsymbol{P}_{aug}$. Note that they are all far away from p in the sense that they do not intersect the edges $p \vee q_i$. Furthermore, the vertices of $\boldsymbol{P}$ are *all* present in $\boldsymbol{P}_{aug}$.

Connector in Dimension $d + 2$

We will construct a $d + 2$-dimensional connector polytope which will have the property that it has four pyramid facets whose (d-dimensional) ground faces are projectively equivalent in every realization. This is the canonical generalization of the construction for $d = 2$ in Chapter 4.

Consider the $d + 1$-dimensional prism $\boldsymbol{P} \times \{0, 1\}$ and q the intersection of the edges connecting the bottom facet $\boldsymbol{P} \times \{0\}$ to the top facet $\boldsymbol{P} \times \{1\}$ (q is on the plane at infinity).

The transmitter polytope is the Lawrence extension of this prism with respect to q. In every realization the two copies of the combinatorial polytope P in $\boldsymbol{prism}(P)$ are projectively equivalent. (Also each realization of P can be completed to a realization of the transmitter polytope.) This is so since $\mathcal{F}_0$ contains all facets except the top and bottom facet of the prism and therefore the existence of a point q conforming with $\mathcal{F}_0$ implies that the edges connecting corresponding vertices of top and bottom faces of the pyramid also meet in q. Hence the bottom face of the pyramid is a projection of the top face.

Along the facet $\boldsymbol{prism}(P)$ we glue another one of these transmitter polytopes and obtain a polytope that prescribes that the four facets $\boldsymbol{pyr}(P)$ are projectively equivalent. Notice that $\boldsymbol{pyr}(P)$ is necessarily flat even though P might not be (we showed in Lemma 4.9 that prisms are always necessarily flat).

Gluing Everything Together

Consider the polytopes $P^{pre}_{p \vee q_i}(\boldsymbol{P}_{aug})$. All of them have a facet which is combinatorial equivalent to $\boldsymbol{P}_{aug}$. (This is the facet which in every realization has the same line image on $p \vee q_i$ as $\boldsymbol{P}_{aug}$.) We glue the pyramids over these polytopes along the pyramids over this facet to connector polytopes (see Figure 5.14).

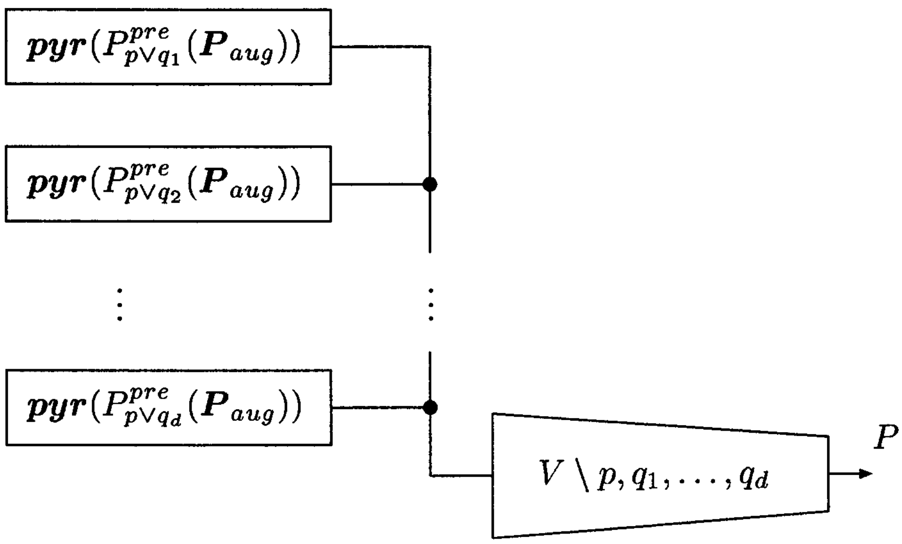

Figure 5.14: *Gluing diagram of the polytope which prescribes* $\boldsymbol{P}$

To the last connector polytope we glue a $d + 2$-dimensional vertex forgetter polytope which forgets the vertices $p, q_1, \ldots, q_d$. This polytope works exactly like the 4-dimensional vertex forgetter. We leave out the details since the construction almost identical to the construction of the connector polytope.

Since all gluing facets are necessarily flat (pyramids), we get a polytope such that in every realization the d-faces which are combinatorially equivalent to $\boldsymbol{P}_{aug}$ have line images on $p \vee q_i$ which are projectively equivalent to the line images of $\boldsymbol{P}_{aug}$. It follows by Lemma 5.11 that these faces must be projectively equivalent to $\boldsymbol{P}_{aug}$. Chopping off the superfluous vertices leaves us with a polytope which is projectively equivalent to $\boldsymbol{P}$. This is exatly the face present at the small end of the vertex forgetter. Hence this polytope prescribes that one face is projectively equivalent to $\boldsymbol{P}$.

Chapter 6

MAXTRIANG∂FL(5) is Hard

This chapter is dedicated to the proof that the problem of finding a maximal number of simplices in a triangulation of the boundary of a d-polytope over all realizations is hard where $d \geq 5$. To give the precise statement we need to set up some terminology: For a polytope $\boldsymbol{P}$ the *boundary* $\partial \boldsymbol{P}$ is the union of all its faces. A *triangulation of the boundary* of a polytope is defined canonically: It is a set of $d-1$-dimensional simplices whose union is the polytope boundary such that any two simplices meet in a (possibly empty) common face.

Consider the following decision problem:

MAXTRIANG∂FL(d)	
Given:	A face lattice L of a d-dimesional polytope and a integer K
Question:	Is there a realization $\boldsymbol{P}$ of L such that there is a triangulation of $\partial \boldsymbol{P}$ with more than K simplices?

This chapter is dedicated to the proof of the following theorem:

Theorem 6.1 *The problem* MAXTRIANG∂FL(d) *for* $d \geq 5$ *is as hard the existential theory of the reals.*

For the definition of the existential theory of the reals see Section 5.1. This hardness implies NP-hardness. Note that it is unclear whether this problem is in NP since no polynomial certificate is known (there are already 4-polytopes which in every realization have an exponential encoding length).

The proof will use the following different realizations of the hexagonal prism: first the upright prism over a regular hexagon, second the upright prism over a slight perturbation of the regular hexagon. If the ground faces are precisely

$$\begin{aligned}
\boldsymbol{G}_{14} &= \operatorname{conv}\left\{\begin{pmatrix}-2\\0\end{pmatrix},\begin{pmatrix}-1\\1\end{pmatrix},\begin{pmatrix}1\\1\end{pmatrix},\begin{pmatrix}2\\0\end{pmatrix},\begin{pmatrix}1\\-1\end{pmatrix},\begin{pmatrix}-1\\-1\end{pmatrix}\right\},\\
\boldsymbol{G}_{17} &= \operatorname{conv}\left\{\begin{pmatrix}-2\\0\end{pmatrix},\begin{pmatrix}-1\\1\end{pmatrix},\begin{pmatrix}1\\1\end{pmatrix},\begin{pmatrix}2\\0\end{pmatrix},\begin{pmatrix}1.5\\-0.5\end{pmatrix},\begin{pmatrix}-1.5\\-0.5\end{pmatrix}\right\},
\end{aligned}$$

then the polytopes $\boldsymbol{P}_i = \boldsymbol{G}_i \times [0,1]$ have maximal triangulations 14 and 17, respectively (see Figure 6.1). These number were computed using DeLoera's UNIVERSALBUILDER [23] and CPLEX.

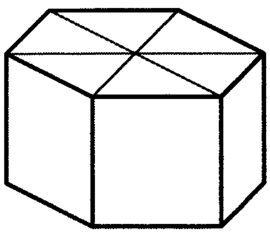

$\boldsymbol{P}_{14}$
maximal 14 tetrahedra

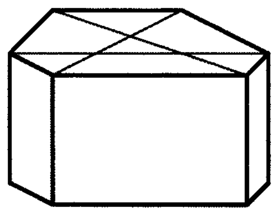

$\boldsymbol{P}_{17}$
maximal 17 tetrahedra

Figure 6.1: *Two realizations of the hexagonal prism with different maximal triangulations*

In order to show that MAXTRIANG∂FL(5) is as hard as the ETR, we will encode an ETR instance S (a system of polynomial inequalities and equations) into a combinatorial polytope. This encoding will be in several steps, each of them will use the following notation:

Notation. A polytope P *prescribes* properties E_0 *or* E_1 of one of its faces F *depending on (the solvability of) the driving system* S if

1. in the case that S has no solution, P prescribes F to have E_0 and
2. in the case that S has solutions, P prescribes F to have E_0 or E_1.

The proof is structured as follows. At the beginning we will work in four dimensions and only in the last step go into the fifth dimension. Given a driving polynomial system S we we will construct a combinatorial polytope which in every realization has a prism over a hexagonal ground face which is projectively equivalent to the hexagons $\boldsymbol{G}_{14}$ or $\boldsymbol{G}_{17}$, depending on S. Following the above notation we mean that $\boldsymbol{G}_{17}$ can only occur in a realization if the system S has a solution; but in any case there is always a realization with $\boldsymbol{G}_{14}$.

By adding more construction elements and lifting the construction to the fifth dimension we get a combinatorial polytope which has a pyramid over $\boldsymbol{P}_{14}$ or $\boldsymbol{P}_{17}$, again depending on the solvability of S.

The construction will be such that all other facets are pyramids over pyramids. A pyramid over a pyramid over a polygon has the same number of simplices in every realization (namely $n-2$ if it is an n-gon). So the maximal triangulation of the boundary of polytope will only depend on the triangulation of the pyramid over the prism which is either $\boldsymbol{P}_{14}$ or $\boldsymbol{P}_{17}$. Since it can be $\boldsymbol{P}_{17}$ only if S had a solution we get that by maximimally triangulating we can decide whether S has a solution.

We have to spend some time checking that the encoding length of the polytope is polynomial in the encoding length of S.

In Section 6.1.2 we will show the start of the construction: Given a driving polynomial system S and two sets $1 < \alpha_1 < \ldots < \alpha_m$ and $1 < \beta_1 < \ldots < \beta_m$ of algebraic numbers we will construct a Shor normal form S_S^{Shor} such that a solution of S_S^{Shor} has the α_i as a solution for the first m variables if and only if S has a solution. The β_i however are always part of a solution of S_S^{Shor} and there are no other partial solutions of S_S^{Shor}.

In Section 6.1.3 we will use this and Richter-Gebert's universality theorem (Theorem 4.13) in order to construct a polytope with a polygonal face G which is prescribed to encode the α_i or the β_i. But it prescribes the α_i only if S has a solution.

Using this construction, the polytope prescribes only the edge slopes of normal polygons. This is not enough: for instance the polygons $\boldsymbol{G}_{14}$ and $\boldsymbol{G}_{17}$ have the same edge slopes. In order to prescribe the exact shape of a polygon we do the same tricks as in the last chapter: We flatten $\boldsymbol{G}_{14}$ and $\boldsymbol{G}_{17}$ until we can double them such that the results are centrally symmetric (convex) polygons. Then by forgetting vertices and prescribing the slopes of these polygons we can prescribe edge and 1-chord slopes. In Section 6.1.4 we supply the details.

On the run we analyse the various proofs and show that all constructions are doable in running time which is polynomial in the encoding length of the driving system.

Up to this point we have constructed a combinatorial 4-polytope which prescribes that a hexagonal prism is prescribed to have the ground face $\boldsymbol{G}_{14}$ or $\boldsymbol{G}_{17}$, depending on the driving system S. Two requirements still have to be met: all other facets must be pyramids and the prisms must not only have the right ground faces, but really be projectively equivalent to $\boldsymbol{P}_{14}$ or $\boldsymbol{P}_{17}$, depending on the driving system. In order to achieve the first, we glue pyramids to all other facets. In Section 6.2.1 we will show that this does not change anything in the prescribing properties of the polytope. This has the effect that all facets (except the hexagonal prism) are pyramids over polygons.

In Section 6.2 we will show what kind of gadget will ensure that the prism is prescribed to be one of $\boldsymbol{P}_{14}$ and $\boldsymbol{P}_{17}$. We will have to lift the polytope to the fifth dimension by considering the bipyramid over it. The facets of this bipyramid are pyramids over the facets of the original polytope, for each original facet there are two pyramids. We end up with a polytope all of whose facets are pyramids over pyramids over polygons except one pyramid over a hexagonal prism which is prescribed to be projectively equivalent to $\boldsymbol{P}_{14}$ or $\boldsymbol{P}_{17}$, depending on the driving system S.

6.1 Prescribing Polygons Depending on a Driving System

Given a polynomial system S we want to construct a combinatorial polytope whose maximal boundary triangulation over all realizations has different sizes, depending on whether this *driving system* S has a solution or not.

6.1.1 Encoding Length

The *encoding length of a polynomial system S with algebraic coefficients* is the sum of the encoding lengths of the participating polynomials. The *encoding length of a polynomial* is the sum of the encoding lengths of each monomial. The *encoding length of a monomial* is the encoding lengths of its coefficients plus the sum of the encoding lengths of its exponents. The *encoding length of an integer* is the ceiling of the dual logarithm of its absolute value plus 1. For instance $17x^4y^3$ has encoding length $\lceil \log_2 17 + 1 \rceil + \lceil \log_2 4 + 1 \rceil + \lceil \log_2 3 + 1 \rceil = 6 + 3 + 3 = 12$. The encoding length of 0 is 0. The *encoding length of a rational number* is the sum of the encoding lengths of its denominator and of its numerator. The *encoding length of an algebraic number α* represented by the irreducible integer polynomial f and the lower and upper bounds l and r is the sum of the encoding lengths of f, l, and r.

We define the *encoding length of a combinatorial polytope* as the product of the number of vertices and the number of facets. (We could define it as the number of vertex-facet incidences, but in fixed dimensions the two definitions are polynomially bounded one by the other.)

6.1.2 A Driving Polynomial System

The first step in our construction is the construction of a Shor normal form such that a partial solution has different values, depending whether the driving system S has a solution or not.

Let S_0 and S_1 be polynomial systems in variables $x_1, \ldots, x_m$ and S be the driving polynomial system in the variables $y_1, \ldots, y_n$. Think of S_1 as the system which determines the coordinates of a polygon with a large triangulation, and S_0 as a system which has a polytope with a small triangulation as a solution. (This will be made clearer below.) We want to construct a polynomial system S' in the variables x_i and y_j such that $x_1, \ldots, x_m$ are always a solution of S_0 or of S_1, but they can only be a solution of S_1 if S has a solution. So if a polytope prescribes for a polygon to encode the solutions of S' then after forgetting all variables except $x_1, \ldots, x_m$ the existence of a solution of S is encoded in the different realizations of this polygon, i.e. the realization can only encode a solution of S_1 if S had a solution.

Let F (F_1 and F_0, respectively) be the set of polynomials f that occur in the equations $f = 0$ of S (S_1 and S_0, resp.) and G (G_1 and G_0, resp.) be the set of polynomials g that occur in the inequalities $g < 0$ of S (S_1 and S_0, resp.). We define the system S' in the variables $x_1, \ldots, x_m, y_1, \ldots, y_n$, and additional variables L and $z_1, \ldots, z_k$ to be

$$\begin{array}{rcll}
f(y_1, \ldots, y_n) \cdot L & = & 0 & \text{for all } f \in F \\
g(y_1, \ldots, y_n) \cdot z_j^2 + L & = & 0 & \text{for all } g \in G \text{ and some unused index } j \\
L \cdot (1 - L) & = & 0 & \\
L \cdot f(x_1, \ldots, x_m) & = & 0 & \text{for all } f \in F_1 \\
L \cdot g(x_1, \ldots, x_m) & < & 1 - L & \text{for all } g \in G_1 \\
(1 - L) \cdot f(x_1, \ldots, x_m) & = & 0 & \text{for all } f \in F_0 \\
(1 - L) \cdot g(x_1, \ldots, x_m) & < & L & \text{for all } g \in G_0
\end{array}$$

Let $V(S) \subseteq \mathbb{R}^n$ be the solution set of S, i.e.

$$\begin{array}{rlll}
V(S) = & \{(y_1, \ldots, y_n) & | & f(y_1, \ldots, y_n) = 0 \text{ for all } f \in F \text{ and} \\
 & & & g(y_1, \ldots, y_n) < 0 \text{ for all } g \in G\}.
\end{array}$$

Similarly, let $V(S_1), V(S_0) \subseteq \mathbb{R}^m$ be the solution sets of S_1 and S_0.

Lemma 6.2 *1. If $V(S) = \emptyset$ then $\pi_x(V(S')) = V(S_0)$ and*

2. if $V(S) \neq \emptyset$ then $\pi_x(V(S')) = V(S_0) \cup V(S_1)$, *where π_x is the projection*

$$\begin{aligned}\mathbb{R}^{n+m+k+1} &\to \mathbb{R}^m,\\ (y_1,\ldots,y_n,x_1,\ldots,x_m,L,z_1,\ldots,z_k) &\mapsto (x_1,\ldots,x_m).\end{aligned}$$

PROOF: In any solution of S', L is either 0 or 1. If we want to have $L = 0$ then the equations and inequalities involving $y_1,\ldots,y_n$ are trivially satisfied with $z_j = 0$. For $x_1,\ldots,x_m$ we only have the obstructions $f(x_1,\ldots,x_m) = 0$ for all $f \in F_0$ and $g(x_1,\ldots,x_m) < 0$ for all $g \in G_0$. If $L = 1$ then $f(y_1,\ldots,y_n) = 0$ for all $f \in F$ and $g(y_1,\ldots,y_n) \cdot z_j^2 = -1$. For this to hold it is necessary and sufficient that $z_j = \sqrt{-1/g(y_1,\ldots,y_n)}$ and $g(y_1,\ldots,y_n) < 0$. Furthermore, the only obstructions for $x_1,\ldots,x_m$ are $f(x_1,\ldots,x_m) = 0$ for all $f \in F_1$ and $g(x_1,\ldots,x_m) < 0$ for all $g \in G_1$. $\square$

Depending on the existence of a solution of the driving system S we want to prescribe different polygon shapes. For this we need Shor normal forms. Also, it will turn out to be sufficient that S_1 and S_0 have only one solution each, namely the edge slopes of given polygons.

Let S_1 be the polynomial system whose only solution for its variables $x_1,\ldots,x_m$ are algebraic numbers $1 < \alpha_1 < \ldots < \alpha_m$. Let similarly S_0 be the polynomial system that defines algebraic numbers $1 < \beta_1 < \ldots < \beta_m$. Then by Lemma 5.2 we can transform S' into a Shor normal form. We denote this $S_S^{\text{Shor}} = S_S^{\text{Shor}}(\alpha_1,\ldots,\alpha_m|\beta_1,\ldots,\beta_m)$

Lemma 6.3 *1. If S has solutions then in any solution of S_S^{Shor} the first m variables are assigned to the α_i or to the β_i.*

2. If S has no solution then in any solution of S_S^{Shor} the first m variables are assigned to the α_i.

3. The system S_S^{Shor} can be computed in time which is polynomial in the encoding length of S. (The α_i and β_i are considered as constants.)

PROOF: The first two parts of the lemma follow directly from Lemma 6.2. We only have to prove the polynomiality of the construction.

The system S' constructed from S, S_1 and S_0 according to Lemma 6.2 clearly has encoding length linear in the encoding length of S. So we only have to check the polynomiality of the construction from Lemma 5.2. For this we follow the construction step by step.

In the first step we decomposed the polynomial equations and inequalities into simple equations and inequalities by introducing (many) auxiliary variables. Note however, that the number of equations and inequalities we used to encode the various intermediate terms is always linear in the encoding length of these terms: this is so for V_0 and V_1, inductively for all V_{2^i}, for all V_α where α is a positive integer and then a negative integer, rational number and algebraic number; the same inductive buildup leads to polynomial encoding length for auxiliary variables for monomials and polynomials.

In the second step we introduce a constant number of variables for each old variable and a constant number of equations and inequality for each old equation and inequality. Since all operations are simple the encoding length of them is also constant. So the encoding length is linear in the enconding length of the system coming out of the first step.

In the third step we introduce new variables for each old variable and each equation and inequality. Also for each operation we introduce a constant number of new operations. Again this means only adding a linear factor to the encoding length. We end up with a Shor normal form S_S^{Shor} which has polynomial (even linear) encoding length in the original system S. $\square$

6.1.3 Prescribing Edge Slopes Depending on Driving System

We have seen that from a driving polynomial system S we can construct a Shor normal form which depending on the existence of a solution to S has one or two predefined partial solutions α_i or β_i for the first m variables. This Shor normal form we can encode in a polytope using Richter-Gebert's universality construction (Theorem 4.13). Using this we will be able to prescribe the slopes of a polygonal face to be fixed (if S has no solution) or to take on one of two states (if S has solutions).

Let $\boldsymbol{G}_1$ and $\boldsymbol{G}_0$ be two normal polygons with algebraic coordinates and $\alpha_1, \ldots, \alpha_m$ and $\beta_1, \ldots, \beta_m$ the numbers they encode as a computation frame (see Section 4.3.1). S is still the driving polynomial system. We define the polytope $P_S^{sl}(\boldsymbol{G}_1|\boldsymbol{G}_0)$ to be

$$P(S_S^{\text{Shor}}(\alpha_1, \ldots, \alpha_m|\beta_1, \ldots, \beta_m))\#(E \setminus x_{m+1}, x'_{m+1}, \ldots, x_n, x'_n),$$

i.e. the polytope which encodes the Shor normal form $S^{\text{Shor}}(\alpha_1, \ldots, \alpha_m|\beta_1, \ldots, \beta_m)$ glued to edge forgetters that forget the edges corresponding to all variables except $x_1, \ldots, x_m$.

Lemma 6.4 *$P^{sl}_S(\mathbf{G}_1|\mathbf{G}_0)$ has a polygonal face G such that*

1. *if S has no solution P prescribes G to have edge slopes projectively equivalent to the edge slopes of $\mathbf{G}_0$, and*

2. *if S has solutions then P prescribes G to have edge slopes projectively equivalent to either the edge slopes of $\mathbf{G}_1$ or to the edge slopes of $\mathbf{G}_1$.*

The encoding length of P is polynomially bounded by the enconding length of S ($\mathbf{G}_i$ are considered as constant).

PROOF: The ground face of the pyramid along which the edge forgetter is glued is prescribed to be a normal computation frame encoding a solution to $S^{\text{Shor}}(\alpha_1, \ldots, \alpha_m|\beta_1, \ldots, \beta_m)$. At the other end of the edge forgetter the edges corresponding to the auxiliary variables have been forgotten, hence the ground face of *that* pyramid is prescribed to be a computation frame encoding the α_i and the β_i, depending on S. But this means that this 2-face has projectively equivalent edge slopes to $\mathbf{G}_1$ or $\mathbf{G}_0$, depending on S. This proves the first two statements. It remains to prove the polynomiality of the construction.

S_S^{Shor} has polynomial encoding length in S by Lemma 6.3. Richter-Gebert's construction of $P(S_S^{\text{Shor}})$ is such that for each simple addition in S_S^{Shor} there is one addition gadget $P^{x+y}[0, 1, x_i, x_j, x_k, \infty]$, an edge-forgetter for the $n - 3$ edge pairs corresponding to the variables which do not occur in this addition and one connector. Also for each simple multiplication there is one multiplication gadget $P^{x \cdot y}[0, 1, x_i, x_j, x_k, \infty]$, an edge forgetter and a connector.

Since the number of auxiliary variables $n - m$ was polynomial in the encoding length of S, this polytope as well as the edge forgetter for the auxiliary variables have an encoding length which is polynomial in the encoding length of S. □

6.1.4 Prescribing Arbitrary Polygons Depending on Driving System

Now we will show how to construct a combinatorial polytope encoding a driving system S which prescribes that a specified 2-face G is projectively equivalent to $\mathbf{G}_{14}$ or to $\mathbf{G}_{17}$, but to $\mathbf{G}_{17}$ only if S has a solution. Also G will be the ground face of a prism facet. The construction is similar to the one in Section 5.2.4.

First we apply a projective transformation to both polygons and obtain "flat" polygons:

$$\begin{aligned} G_{14} &\sim \operatorname{conv}\left\{\begin{pmatrix}-3\\0\end{pmatrix},\begin{pmatrix}-1\\-1\end{pmatrix},\begin{pmatrix}1\\-1\end{pmatrix},\begin{pmatrix}7/2\\0\end{pmatrix},\begin{pmatrix}17/3\\13/3\end{pmatrix},\begin{pmatrix}-3\\13/3\end{pmatrix}\right\}, \\ G_{17} &\sim \operatorname{conv}\left\{\begin{pmatrix}-3\\0\end{pmatrix},\begin{pmatrix}-1\\-1\end{pmatrix},\begin{pmatrix}1\\-1\end{pmatrix},\begin{pmatrix}7/2\\0\end{pmatrix},\begin{pmatrix}45/11\\13/11\end{pmatrix},\begin{pmatrix}-3\\13/11\end{pmatrix}\right\}, \end{aligned}$$

These we double as in Section 5.2.6 and get centrally symmetric, hence nor-

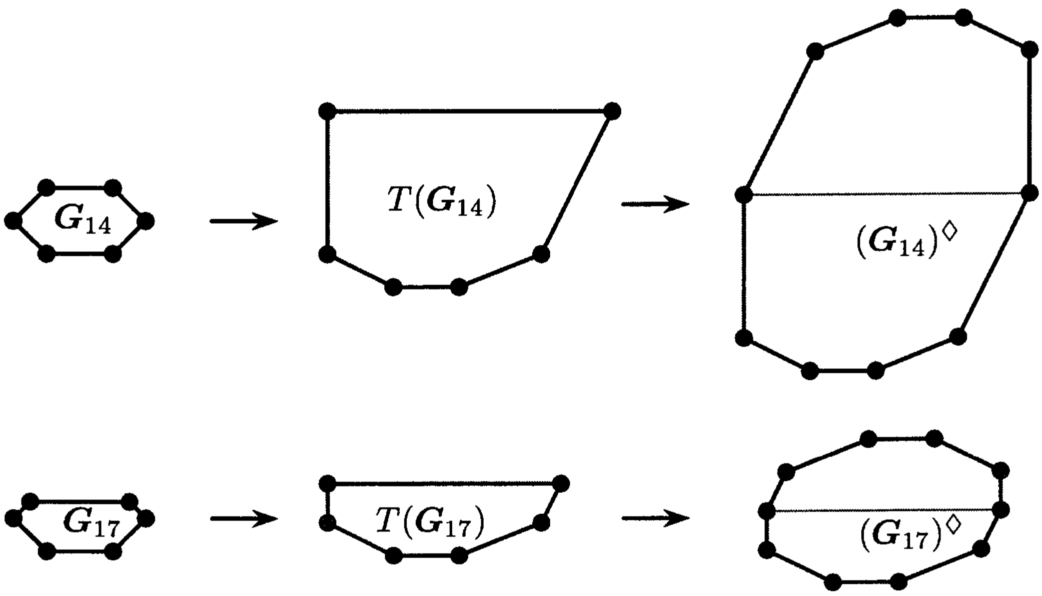

Figure 6.2: *Constructions of* $(G_{14})^\diamond$ *and* $(G_{17})^\diamond$

mal, polygons $(G_{14})^\diamond$ and $(G_{17})^\diamond$ (see Figure 6.2). The slopes of these polygons are:

	$s(1)$	$s(1,2)$	$s(2)$	$s(2,3)$	$s(3)$	$s(3,4)$	$s(4)$	$s(4,5)$	$s(5)$
$(G_{14})^\diamond$	∞	$-\frac{8}{3}$	$-\frac{1}{2}$	$-\frac{1}{4}$	0	$\frac{2}{9}$	$\frac{2}{5}$	$\frac{8}{7}$	2
$(G_{17})^\diamond$	∞	$-\frac{12}{11}$	$-\frac{1}{2}$	$-\frac{1}{4}$	0	$\frac{2}{9}$	$\frac{2}{5}$	$\frac{12}{17}$	2

As in Section 5.2.4 we will one by one forget the vertices $1 \wedge 2$, $2 \wedge 3$, $3 \wedge 4$, and $4 \wedge 5$ as well as their opposite vertices and prescribe the slopes of these polytopes.

Note that all edge slopes and *all but two* 1-chord slopes are identical. Hence when we delete $2 \wedge 3$, the resulting polytopes $(\boldsymbol{G}_{14})_2^{\diamond}$ and $(\boldsymbol{G}_{17})_2^{\diamond}$ have the same slopes. So we can prescribe these common slopes with the polytope $P^{sl}((\boldsymbol{G}_{14})_2^{\diamond})$. (This polytope does *not* depend on the driving system S.) The same we can do for the deletion of vertex $3 \wedge 4$.

We use the polytopes $P_S^{sl}((\boldsymbol{G}_{14})_1^{\diamond}|(\boldsymbol{G}_{17})_1^{\diamond})$ and $P_S^{sl}((\boldsymbol{G}_{14})_4^{\diamond}|(\boldsymbol{G}_{17})_4^{\diamond})$ prescribe slopes depending on the driving system S. To the last connector we glue a vertex forgetter which forgets the auxiliary vertices. See the upper part of the gluing diagram in Figure 6.4.

Consider a realization of this polytope. The polygon $sl(\boldsymbol{G}^{\diamond})$ which is the common polygon of the connectors on the right-hand side of the diagram has the same edge slopes as $(\boldsymbol{G}_{14})^{\diamond}$ and $(\boldsymbol{G}_{17})^{\diamond}$. It also has the same 1-chord slopes $s(2,3)$ and $s(3,4)$ since this is prescribed by $P^{sl}((\boldsymbol{G}_{14})_2^{\diamond})$ and $P^{sl}((\boldsymbol{G}_{14})_3^{\diamond})$. Both $P_S^{sl}((\boldsymbol{G}_{14})_1^{\diamond}|(\boldsymbol{G}_{17})_1^{\diamond})$ and $P_S^{sl}((\boldsymbol{G}_{14})_4^{\diamond}|(\boldsymbol{G}_{17})_4^{\diamond})$ encode the driving system and depending on the existence of a solution, the 1-chord slopes of $s(1,2)$ and $s(4,5)$ can be the ones of $\boldsymbol{G}_{14}$ or $\boldsymbol{G}_{17}$.

The trouble is if our driving system has a solution, then the two polygons $sl(\boldsymbol{G}_1^{\diamond})$ and $sl(\boldsymbol{G}_4^{\diamond})$ can have inconsistent slopes for $s(1,2)$ and $s(4,5)$, one coming from $\boldsymbol{G}_{14}$ and one from $\boldsymbol{G}_{17}$. By Lemma 5.7 each of the four combinations gives a polygon which is determined up to projective equivalence: we have a 10-gon where all edge slopes and all 1-chord slopes except $s(5,1')$ and $s(5',1)$ are known. Figure 6.3 shows the four possible common realizations where the auxiliary vertices $1' \wedge 2', \ldots, 4' \wedge 5'$ have already been forgotten.

But notice that among these four polygons only $\boldsymbol{G}_{14}$ and $\boldsymbol{G}_{17}$ are normal. So we glue another polytope to a connector, a *normalizer*. This is a polytope which only ensures that its gluing pyramid has a normal ground face. We can obtain such a polytope can for instance by encoding an empty Shor normal form (i.e. no variables). Or we can use Richter-Gebert's $H(1,2,3,4,1',2',3',4')$ polytope whose gluing facet is a pyramid over a normal octagon with the *only* restriction that the edge slopes are harmonic, i.e. $cr(s(1),s(2),s(3),s(4)) = -1$. Then we can just forget edges 4 and $4'$ and get a necessarily normal hexagon (see Figure 6.5 which has no more restrictions. So among the four mentioned realizations only the desired two, $\boldsymbol{G}_{14}$ and $\boldsymbol{G}_{17}$, are eventually possible.

To the last normalizer we glue a transmitter (not a connector, just half of it). Call the resulting polytope $P_S(\boldsymbol{G}_{14}|\boldsymbol{G}_{17})$. We conclude:

Lemma 6.5 *The polytope $P = P_S(\boldsymbol{G}_{14}|\boldsymbol{G}_{17})$ has a face F which is a hexagonal prism.*

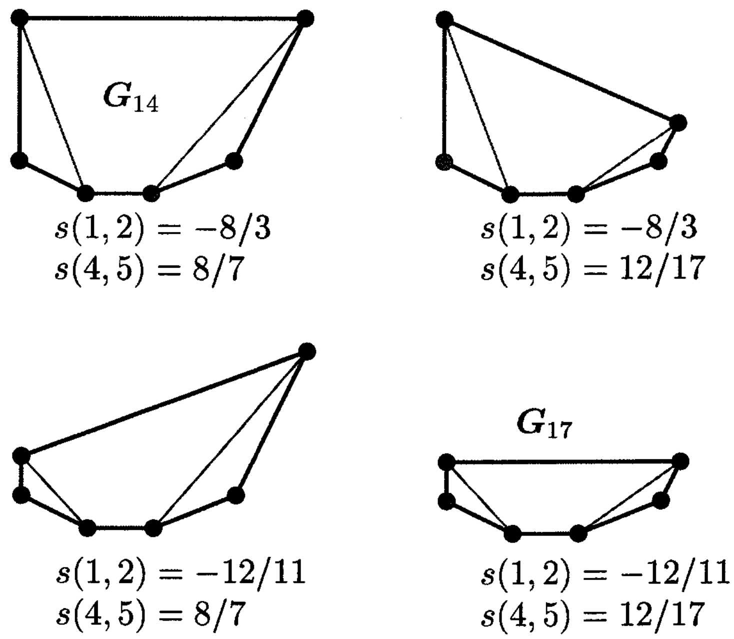

Figure 6.3: *The four possible realizations before normalizing*

1. *If S has a solution then F is prescribed to have top and bottom faces which are projections of each other and which are projectively equivalent to $\boldsymbol{G}_{14}$ or $\boldsymbol{G}_{17}$.*

2. *If S has no solution then F is prescribed to have top and bottom faces which are projections of each other and which are projectively equivalent to $\boldsymbol{G}_{14}$.*

The encoding length of P is polynomial in the encoding length of S.

Are we done yet? We have constructed a 4-polytope that prescribes that one of its a 3-prism facets has both top and bottom polygon projectively equivalent to $\boldsymbol{G}_{14}$ or $\boldsymbol{G}_{17}$, depending on the existence of a solution of a driving polynomial system. Notice that this prism is otherwise free to take any shape and in particular that this does not mean that the prism is prescribed to be projectively equivalent to $\boldsymbol{P}_{14}$ (or $\boldsymbol{P}_{17}$): the bottom and top face might not be parallel. In that case the maximal triangulations would be different. Check out the prism in Figure 6.6. The ground face is a regular 6-gon $\boldsymbol{G}_{14}$. The side edges are all parallel, but the vertices in the top face have heights 1 (front edge), 3 (next two vertices), and 5 (back edge). This polytope has a maximal triangulation of even 18 tetrahedra (so it is neither $\boldsymbol{P}_{14}$ nor $\boldsymbol{P}_{17}$).

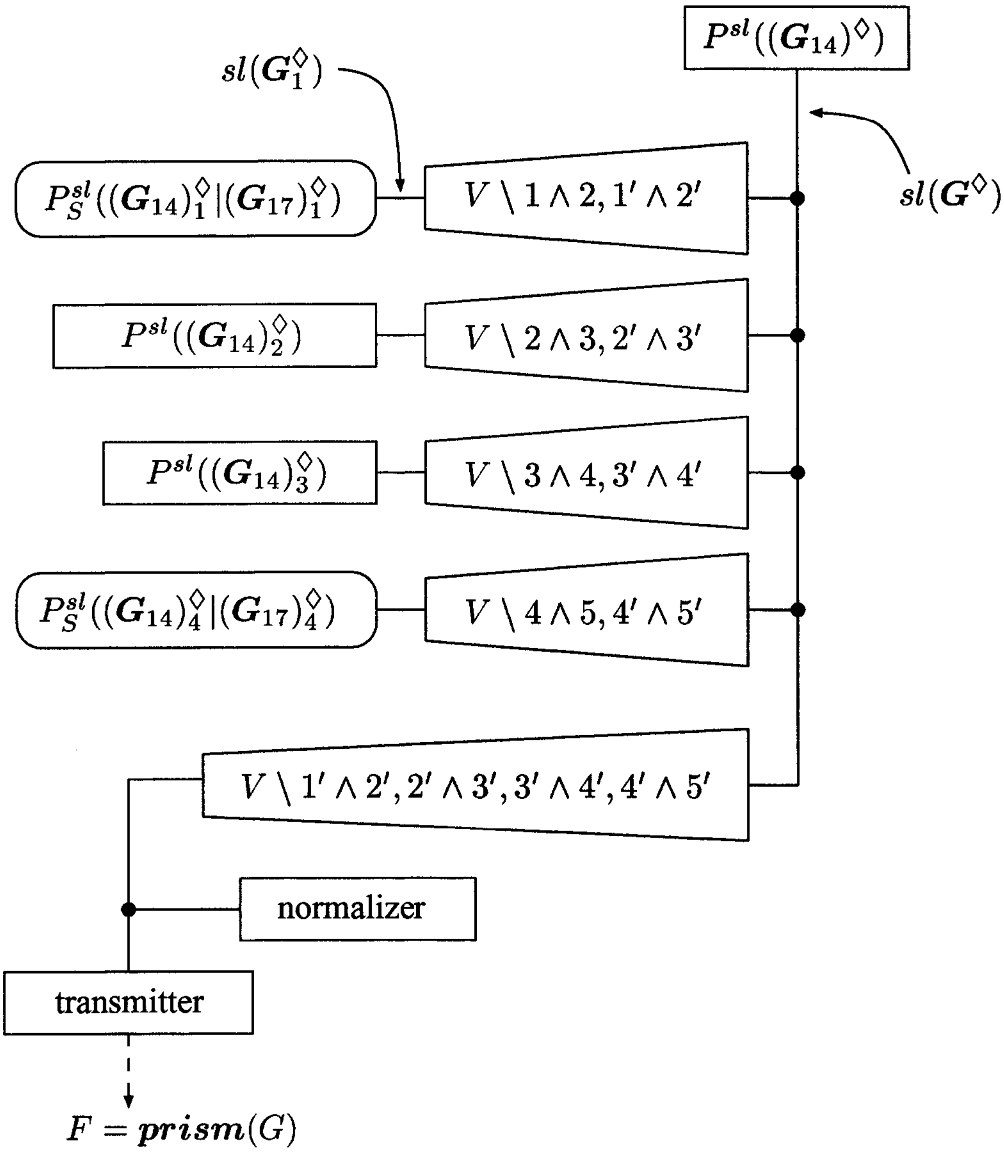

Figure 6.4: *Gluing diagram of the polytope* $P_S(G_{14}|G_{17})$

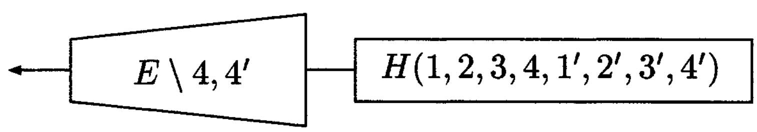

Figure 6.5: *Gluing diagram of the normalizer*

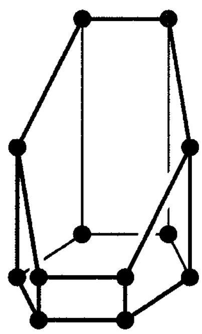

Figure 6.6: *A prism with ground face G_{14}, but a maximal triangulation with* 18 *tetrahedra. Nota: DeLoera, Santos, and Takeuchi [25, 58] proved that this realization has the maximal triangulation of the hexagonal prism.*

6.2 Prescribing Hexagonal Prisms Depending on Driving System

We have seen now how to construct a combinatorial polytope which encodes a driving polynomial system in a way that one facet is a prism over a special polygon G_{14} or G_{17}, but G_{17} is only possible if the system had a solution. We promised, however, to construct a polytope

1. such that one facet must be projectively equivalent to one of the hexagonal prisms $P_{14} = G_{14} \times [0,1]$ or $P_{17} = G_{17} \times [0,1]$ — depending on the existence of a solution for S — and
2. such that all other facets are pyramids over polygons.

We first tackle the second problem, then the first.

6.2.1 All Other Facets are Easy to Triangulate

For all facets F of $P = P_S(G_{14}|G_{17})$ which are neither pyramids nor the special hexagonal prism facet which we have prescribed we do the following operation: Consider the pyramid $\boldsymbol{pyr}(F)$, this has the ground facet F. Along this face we glue P, i.e. $P' = P\#_F\boldsymbol{pyr}(F)$. We do this for all facets which are not pyramids, one by one. The result $P^{pyr} = P_S^{pyr}(G_{14}|G_{17})$ is a polytope which has only pyramid facets (except the special hexagonal prism), but does it still prescribe anything?

Yes, it does since all facets F were necessarily flat, so in each realization of P^{pyr} we could chop the pyramids back off and obtain a realization of P. Conversely, given a realization of P we can place a point *beyond* each non-pyramid facet (except the special hexagonal prism), take the convex hull and obtain a realization of P^{pyr}. The reason that all facets are necessarily flat is that in our construction we only used the building blocks of Richter-Gebert's construction, and those are the transmitters, edge forgetters, the polytope X, the pyramid over a tent over an octagon, and a polytope which we will see in a minute (Richter-Gebert calls it a "slope transmitter"). It is easy to check that all faces of these are indeed of the three categories.

6.2.2 The Parallelifier

In this section we will only talk about G_{14}. However, our argument eventually only use the combinatorial polytope, hence it all is also valid for G_{17}.

Suppose Q is a hexagonal prism whose bottom face $G = G(1, \dots, 6)$ and top face $G' = G(1', \dots, 6')$ are projections of each other and which are in turn projectively equivalent to G_{14}. Since G_{14} is normal, so are G and G'. Hence there are lines l_∞ and l'_∞ which include the vertices $1 \wedge 4$, $2 \wedge 5$, and $3 \wedge 6$ ($1' \wedge 4'$, $2' \wedge 5'$, and $3' \wedge 6'$, respectively). We can prescribe a polytope to have these properties.

If we prescribe further that $1 \wedge 4 = 1' \wedge 4'$ and $2 \wedge 5 = 2' \wedge 5'$ then necessarily $l_\infty = l'_\infty$ (they have two points in common). But not only that:

Lemma 6.6 *Q is projectively equivalent to P_{14}.*

PROOF: In lieu of a proof will construct the projective transformation. Since top and bottom faces are projections of each other there is a point q (outside of Q) which is the intersection of edges connecting the top and bottom face. There is an (invertible) projective transformation taking the ground face G to $G_{14} \times \{0\}$. By Lemma 4.15 there is a unique projective transformation f taking q to the intersection of the upright edges of P_{14} and which takes one point of the top face G' to the corresponding vertex of $G_{14} \times \{1\}$.

Then f must map the line l_∞ of G to the corresponding line l^{14}_{infty} of G_{14}. This line is actually at infinity, since opposite edges are parallel in G_{14}. But since $l_\infty = l'_\infty$, the plane $\overline{G'}$ is mapped to a plane which is parallel to $G_{14} \times \{0\}$. This plane must be at $\mathbb{R}^2 \times \{1\}$ since one point in it (the vertex of G') is mapped to its correspondent in $G_{14} \times \{1\}$.

Also the images of the upright edges of Q are all parallel since q is mapped to the intersection points of the upright edges of P_{14} which is at infinity. So these lines map to $\{v\} \times \mathbb{R}$ where v is a vertex of G_{14}. It follows that the vertices of the top face G' are mapped to their correspondents in $G_{14} \times \{1\}$: they are intersections of $\{v\} \times \mathbb{R}$ and $\mathbb{R}^2 \times \{1\}$. The preimage of f of the hyperplane at infinity misses Q since it is the plane spanned by l_∞ and q. It follows that $f(Q) = P_{14}$. □

In order to prescribe that $1 \wedge 4 = 1' \wedge 4'$ we use Lawrence construction again. Start with P_{14}. The parallelifier is the (combinatorial) Lawrence extension $\Lambda(P_{14}, 1 \wedge 4)$. The set $\mathcal{F}^0$ consists of the top and bottom polygon as well as the quadrangular faces $1 \vee 1'$ and $4 \vee 4'$, $\mathcal{F}^+$ are the quadrangular faces $2 \vee 2'$ and $3 \vee 3'$ and $\mathcal{F}^-$ are the quadrangular faces $5 \vee 5'$ and $6 \vee 6'$. See Figure 6.7. This polytope prescribes exactly that $1 \wedge 4 = 1' \wedge 4'$, and we could glue it to

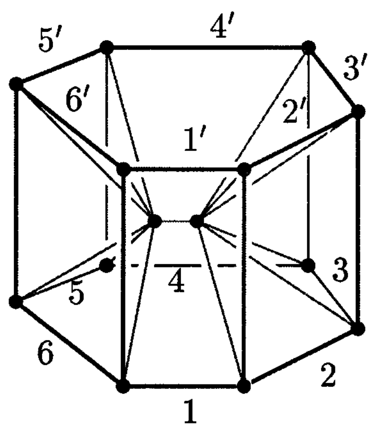

Figure 6.7: *The parallelifier: quadruple of edges must go through a point*

our construction.

Again, as in Section 6.2.1, we can make all non-pyramid facets of this polytope into pyramids by gluing pyramids over them. This does not change the prescribing power since these facets are necessarily flat (pyramids or tents). The resulting polytope we call the *modified parallelifier* $Par_{1,4}$.

The choice of opposite edges is of course not limited to 1 and 4: By relabeling we can construct the polytope $Par_{2,5}$ which prescribe that $2 \wedge 5 = 2' \wedge 5'$. In the next section we show how to attach these polytopes to the polytope prescribing the ground face of the hexagonal prism.

6.2.3 Bipyramids

If we glue a parallelifier to the polytope $P = P_S^{pyr}(G_{14}|G_{17})$ along a hexagonal prism, then this prism is gone—to the interior of the polytope—and it does us no good to have prescribed it. Also, one quadruple would not be enough to really have $G_{14} \times [0,1]$. The way out of this dilemma is to lift the construction to the fifth dimension. First construct a bipyramid over the whole polytope we constructed so far. The bipyramid is defined as $\boldsymbol{bipyr}(P) = \boldsymbol{pyr}(P, p_1)\#_P\boldsymbol{pyr}(P, p_2)$. This polytope $P_S^{bipyr}(G_{14}|G_{17})$ has two new vertices, p_1 and p_2. It is still a combinatorial polytope since it is realizable (with G_{14}). Its facets are exclusively pyramids over the facets of P; two pyramids, one with apex p_1, the other with apex p_2 share a 3-face which is a facet of P. So this bipyramid has only facets which are pyramids over pyramids over polygons and two facets which are pyramids over a hexagonal prism. Note that these facets *share* the hexagonal prism. The ground face of this hexagonal prism is prescribed to be projectively equivalent to G_{14} or G_{17}, depending on the driving system S.

We do not know whether P is necessarily flat. But the bipyramid still prescribes that the 3-face corresponding to the prism facet of P has top and bottom faces projectively equivalent to G_{14} (or G_{17}): All 3-faces of P are still present, in each realization of the bipyramid. All basic building blocks are necessarily flat since they are all Lawrence extensions where $\mathcal{F}^0$ is non-empty (necessarily flat by Lemma 4.8), except the pyramid over a tent which is a pyramid (necessarily flat by Lemma 4.9). Hence each building block prescribes what it is supposed to prescribe and glued together they play together just like before.

By the same token we can construct the bipyramid $Par_{1,4}^{bipyr}$ over the modified parallelifier $Par_{1,4}$. It has many facets which are pyramids over pyramids over polygons (quadrilaterals or hexagons) and two pyramids over a hexagonal prism. This hexagonal prism is prescribed to have $1 \wedge 4 = 1' \wedge 4'$: All 2-faces of the modified parallelifier are there, hence all 3-faces of the parallelifier (which are necessarily flat) lie in 3-dimensional hyperplanes, therefore all vertices coming from the parallelifier (which in turn is necessarily flat) lie in a 4-dimensional hyperplane, so it still has its power to prescribe.

We can glue one pyramid over the hexagonal prism of $P_S^{bipyr}(G_{14}|G_{17})$ to the pyramid over the hexagonal prism of $Par_{1,4}^{bipyr}$. To the other pyramid over the hexagonal prism of $P_S^{bipyr}(G_{14}|G_{17})$ we can glue $Par_{2,5}^{bipyr}$. Let us analyse the facets. There are only pyramids over pyramids over polygons, except two pyramids over a hexagonal prism. Notice that by the gluing construction this is again the *same* hexagonal prism. Since the pyramid over the hexagonal prism (the gluing facet) is necessarily flat, the hexagonal prism is prescribed to be projectively equivalent to $\boldsymbol{P}_{14}$ or $\boldsymbol{P}_{17}$.

6.2.4 Triangulating the Polytope

Note that all triangulations of a pyramid are induced by triangulations of the ground face.

In any realization of this polytope the maximal boundary triangulation only depends on the triangulation of the pyramids over the hexagonal prism: All other facets have only triangulations which are present in any dimension ($n-2$ if they are pyramids over pyramids over n-gons). Say we have a realization where the hexagonal prism is projectively equivalent to $\boldsymbol{P}_{14}$. Then the two pyramids over it have (each separately) a maximal triangulation of 14 4-simplices. But this maximal triangulation can also be achieved: start with a triangulation of the hexagonal prism with 14 3-simplices. Also triangulate all polygons in any way (except the boundary of the hexagonal prism which is already triangulated).

Now these triangulations can be extended to triangulations of the pyramids over these faces and to the pyramids over the pyramids such that the whole set of 4-simplices is a triangulation: It is easy to see that this forms a subdivision of all facets into 4-simplices. The only thing that could go wrong is that two facets have a different triangulation in a common 3-face. But this 3-face must be either a tetrahedron, a pyramid over a polygon, or the hexagonal prism. But these either have a unique triangulation or we had determined the triangulation, so the two adjacent facets could not induce different triangulations.

The same construction works if the hexagonal prism is projectively equivalent to $\boldsymbol{P}_{17}$. Only then there are 6 4-simplices more. This surplus can only be achieved if S has a solution. It follows that by triangulating the boundary maximally over all realizations we can decide whether S has a solution.

Bibliography

[1] F. Aurenhammer and Y. Xu. Optimal triangulations. Technical Report SFB F003-099, Technische Univ. Graz, Graz, Austria, 1999.

[2] D. Avis and H. ElGindy. Triangulating point sets in space. *Discrete Comput. Geom.*, 2:99–111, 1987.

[3] D. Barnette and B. Grünbaum. Preassigning the shape of a face. *Pacific J. Math.*, 32:299–306, 1970.

[4] M. M. Bayer. Face numbers and subdivisions of convex polytopes. In T. B. et al, editor, *Polytopes: Abstract, Convex and Computational.* Kluwer Academic Publishers, Netherlands, 1994.

[5] A. Below, U. Brehm, J. A. De Loera, and J. Richter-Gebert. Minimal simplicial dissections and triangulations of convex 3-polytopes. *Discrete Comput. Geom.*, 24:35–46, 2000.

[6] A. Below, J. A. De Loera, and J. Richter-Gebert. The complexity of finding small triangulations of convex 3-polytopes. To appear.

[7] A. Below, J. A. De Loera, and J. Richter-Gebert. Finding minimal triangulations of convex 3-polytopes is NP-hard. In *Proc. 11th Symp. Discrete Alg.*, pages 65–66, New York, 2000. ACM. Extended abstract.

[8] M. W. Bern, H. Edelsbrunner, D. Eppstein, S. A. Mitchell, and T.-S. Tan. Edge insertion for optimal triangulation. *Discrete Comp. Geom.*, 10(1):47–65, 1993.

[9] M. W. Bern and D. Eppstein. Mesh generation and optimal triangulation. In D.-Z. Du and F. K.-M. Hwang, editors, *Computing in Euclidean Geometry*, number 4 in Lecture Notes Series on Computing, pages 47–123. World Scientific, 2nd edition, 1995.

[10] M. W. Bern, D. Eppstein, and J. Gilbert. Provably good mesh generation. *J. Comput. System Sci.*, 48(3):384–409, 1994.

[11] L. Billera and A. Björner. Face numbers of polytopes and complexes. In J. E. Goodmann and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 291–310. CRC Press, New York, 1997.

[12] A. Björner, M. L. Vergnas, B. Sturmfels, N. White, and G. M. Ziegler. *Oriented Matroids*. Cambridge University Press, Cambridge, 1992.

[13] J. Bochnak, M. Coste, and M.-F. Roy. *Real Algebraic Geometry*. Ergebnisse der Mathematik und ihrer Grenzgebiete; Folge 3, Vol. 36. Springer-Verlag, Berlin etc., 1998.

[14] J. Böhm. private communication.

[15] J. Canny. Some algebraic and geometric computations in $PSPACE$. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, pages 460–467, 1988.

[16] B. Chazelle. Convex partitions of polyhedra: a lower bound and worst-case optimal algorithm. *SIAM J. Comput.*, 13:488–507, 1984.

[17] B. Chazelle and L. Palios. Triangulating a nonconvex polytope. *Discrete Comp. Geom.*, 5:505–526, 1990.

[18] F. Y. L. Chin, S. P. Y. Fung, and C. A. Wang. Approximation for minimum triangulations of simplicial convex 3-polytopes. *Discrete Comp. Geom.*, 26(4):499–511, Dec 2001.

[19] R. W. Cottle. Minimal triangulation of the 4-cube. *Discrete Math.*, 40:25–29, 1982.

[20] H. T. Croft, K. J. Falconer, and R. K. Guy. *Unsolved Problems in Geometry*. Springer-Verlag, New York, 1991.

[21] J. A. De Loera. Computing minimal and maximal triangulations of convex polytopes. manuscript available at `http://www.inf.ethz.ch/personal/deloera`.

[22] J. A. De Loera, S. Hosten, and F. Santos. The polytope of all trianguations of a point configuration. *Documenta Mathematica*, 1:103–119, 1996.

[23] J. A. De Loera and S. Nota. UNIVERSALBUILDER, a c++ program that generates CPLEX readable integer programs to compute minimal and maximal triangulations of convex polytopes of arbitrary dimension. available on `http://www.math.udavis.edu/deloera`.

[24] J. A. De Loera, J. Rambau, and F. Santos. *Spaces of Triangulations*. to appear.

[25] J. A. De Loera, F. Santos, and F. Takeuchi. Extremal properties for dissections of convex 3-polytopes. *SIAM J. Discrete Math.*, 14(2):143–161, 1999.

[26] M. E. Dyer and A. M. Frieze. On the complexity of computing the volume of a polyhedron. *SIAM J. Comput.*, 17(5):967–974, 1988.

[27] M. E. Dyer, A. M. Frieze, and R. Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *J. Assoc. Comput. Mach.*, 38(1):1–17, 1991.

[28] H. Edelsbrunner, F. P. Preparata, and D. B. West. Tetrahedrizing point sets in three dimensions. *J. Symb. Comp.*, 10:335–347, 1990.

[29] E. El-Mallah and C. J. Colbourn. On two dual classes of planar graphs. *Discrete Math.*, 80:21–40, 1990.

[30] S. Fortune. Voronoi diagrams and Delaunay triangulations. In J. E. Goodman and R. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 377–388. CRC Press, New York, 1997.

[31] M. R. Garey and S. J. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, San Francisco, 1979.

[32] P. Gritzmann and V. Klee. On the complexity of some basic problems in computational convexity. In T. Bisztriczky et al., editors, *Polytopes: Abstract, Convex and Computational*, pages 374–466. Kluwer Academic Publishers, Netherlands, 1994.

[33] M. Haiman. A simple and realatively efficient triangulation of the n-cube. *Discrete Comput. Geom.*, 6:287–289, 1991.

[34] R. B. Hughes and M. R. Anderson. Simplexity of the cube. *Discrete Math.*, 158(1–3):99–150, 1996.

[35] I. Itenberg. Viro's method and t-curves. In *Algorithms in algebraic geometry and applications*, volume 143 of *Progr. Math.*, pages 177–192, Santanter, 1994. Birkhäuser, Basel, 1996.

[36] N. Jacobson. *Basic Algebra*, volume 1. Freeman and Co., New York, 2nd edition, 1985.

[37] P. Kleinschmidt. On facets with non-arbitrary shapes. *Pacific J. Math.*, 65(1):97–101, 1976.

[38] C. W. Lee. Regular triangulations of convex polytopes. In P. Gritzmann and B. Sturmfels, editors, *Applied Geometry and Discrete Mathematics - The Viktor Klee Festschrift*, volume 4 of *DIMACS series in Discrete Math. and Theor. Comp. Sci.*, pages 443–456. 1991.

[39] C. W. Lee. Subdivisions and triangulations of polytopes. In J. E. Goodman and R. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 271–290. CRC Press, New York, 1997.

[40] B. Mishra. Computational real algebraic geometry. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 537–556. CRC Press, New York, 1997.

[41] D. Orden and F. Santos. Asymptotically efficient triangulations of the d-cube. Preprint available at `http://arXiv.org/abs/math.CO/0204157`, 2002.

[42] J. O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford Univ. Press, New York/Oxford, 1987.

[43] L. Palios. Optimal tetrahedralization of the $3d$-region "between" a convex polyhedron and a convex polygon. *Computational Geometry Theory and Applications*, 6(5):263–276, Sep 1996.

[44] J. Rambau. *Polyhedral Subdivisions and Projections of Polytopes*. PhD thesis, TU Berlin, Berlin, 1996.

[45] J. Richter-Gebert. Finding small triangulations of polytope boundaries is hard. to appear in Discrete Comp. Geom., manuscript available at `http://www.inf.ethz.ch/personal/richter`.

[46] J. Richter-Gebert. *Realization Spaces of Polytopes*. Lecture Notes in Mathematics 1643. Springer Verlag, Berlin/Heidelberg/New York, 1996.

[47] N. Robertson and P. D. Seymour. Graph minors—a survey. In I. Anderson, editor, *Surveys in Combinatorics*, page nota. Cambridge University Press, 1985.

[48] A. Robinson. *Complete Theories*. Studies in logic and the foundations of mathematics. North-Holland, Amsterdam, New York, Oxford, 2nd edition, 1997.

[49] G. L. Rothschild and E. G. Straus. On triangulations of the convex hull of n points. *Combinatorica*, 5:167–179, 1985.

[50] M.-F. Roy. The role of hilbert problems in real algebraic geometry. In *Proceedings of the 9th general meeting*, European Women in Mathematics, pages 189–200, Hindawi, 2000.

[51] J. Ruppert and R. Seidel. On the difficulty of triangulating three-dimensional nonconvex polyhedra. *Discrete Comp. Geom.*, 7:227–253, 1992.

[52] E. Schönhardt. Über die Zerlegungen von Dreieckspolyedern in Tetraeder. *Math. Ann.*, 98:309–312, 1928.

[53] A. Seidenberg. A new decision method for elementary algebra. *Ann. Math.*, 60(2):365–374, 1954.

[54] P. Shor. Stretchability of pseudolines is NP-hard. In P. Gritzmann and B. Sturmfels, editors, *Applied Geometry and Discrete Mathematics – The Victor Klee Festschrift*, volume 4 of *DIMACS Ser. Appl. Math. and Theor. Comp. Sci.*, pages 531–545. Amer. Math. Soc., Providence, RI, 1991.

[55] D. D. Sleator, R. E. Tarjan, and W. P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1:647–681, 1988.

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

[57] B. Sturmfels. Polynomial equations and convex polytopes. *Amer. Math. Monthly*, 105:907–922, 1998.

[58] F. Takeuchi. *Combinatorics of Triangulations*. PhD thesis, University of Tokyo, Tokyo, 2001.

[59] M. J. Todd. *The Computation of Fixed Points and Applications*, volume 124 of *Lecture Notes in Econom. and Math. Systems*. Springer-Verlag, Berlin, 1976.

[60] C. A. Wang and B. Yang. Minimal tetrahedralizations of a class of polyhedra. In *Proc. 12th Canad. Conf. Computational Geometry*, Aug 2000.

[61] G. M. Ziegler. *Lectures on Polytopes*. Springer-Verlag, New York, 1994.

Index

Curriculum Vitae

Alexander Below

born on June 19, 1974 in Berlin, Germany

1988–1992	**Specialist Highschool "Heinrich Hertz"** **for Mathematics and Natural Sciences** **in Berlin, Germany** Degree: Abitur
1992–1995	**studies at Technical University Berlin, Germany** major: Mathematics minor: Business Administration and Computer Science Degree: Vordiplom
1995–1997	**studies at Cornell University, Ithaca, N.Y., U.S.A.** Operations Research Degree: Master of Science
1997–2002	**Ph.D. student at ETH Zürich, Switzerland** Theoretical Computer Science