
MINIMAL TRIANGULATIONS AND NORMAL SURFACES

Benjamin Andrew Burton

Submitted in total fulfilment of the
requirements of the degree of
Doctor of Philosophy



THE UNIVERSITY OF
MELBOURNE

Department of Mathematics and Statistics
The University of Melbourne
Australia

September 9, 2003

Abstract

This thesis examines three distinct problems relating to the combinatorial structures of minimal 3-manifold triangulations and to the study of normal surfaces within these triangulations. These problems include the formation and analysis of a census of 3-manifold triangulations, a study of splitting surfaces within 3-manifold triangulations and an investigation into the complexity of the normal surface solution space.

An algorithm for generating a census of all closed prime minimal 3-manifold triangulations is presented, extending the algorithms of earlier authors in several ways. Automorphisms of face pairings are utilised to improve the efficiency of the generation of triangulations. 0-efficiency tests and searches for particular subcomplexes within a triangulation are introduced to aid the subsequent processing of these triangulations. Results involving face pairing graphs are proven for the purpose of eliminating large classes of triangulations at different stages of the algorithm.

Using this algorithm, a census is formed of all closed prime minimal triangulations containing at most six tetrahedra. The census of non-orientable triangulations in particular is the first such census to be published. A detailed analysis is performed of the underlying combinatorial structures of the resulting triangulations, extending current knowledge in both the orientable and non-orientable cases. Also included is a full listing of the vertex normal surfaces of these triangulations.

An infrastructure is then developed for studying splitting surfaces within 3-manifold triangulations. Splitting surfaces represent a particular class of normal surfaces containing only quadrilateral discs, and have several interesting combinatorial and topological properties. Splitting surface signatures are introduced to assist with representation and computation, and a census of all splitting surface signatures of order ≤ 8 is presented.

The final problem, which explores the complexity of the normal surface solution space, concentrates on bounding the number of maximal embedded faces of the projective solution space. An extensive analysis is performed upon the geometric structure of this solution space when represented using edge weight coordinates. In the case of general position this geometric structure is used to formulate a corresponding algebraic problem from which partial results are obtained.

Both the study of the census algorithm and the analysis of the resulting minimal triangulations draw upon previously published results. These results are enhanced and significant new related results are offered. In addition, the computer software written to support this work provides a powerful tool for the study of normal surfaces and 3-manifold triangulations.

The investigations into splitting surfaces and the complexity of the normal surface solution space are entirely new. The theory that is developed constitutes a strong basis for continued research into these areas.

Declaration

This is to certify that

- (i) the thesis comprises only my original work towards the PhD,
- (ii) due acknowledgement has been made in the text to all other material used,
- (iii) the thesis is less than 100,000 words in length, exclusive of tables, maps, bibliographies and appendices.

Benjamin Andrew Burton

Acknowledgements

My foremost thanks are to my supervisor, Prof. J. Hyam Rubinstein, for his continual encouragement and support, for many inspirational conversations, for always finding time to talk and for his remarkable patience when I was otherwise occupied. Thanks also to Prof. William Jaco for making me feel at home during my stays in Oklahoma, for his steadfast encouragement, for his willingness to discuss new ideas and for his remarkable wines. Thanks to Linda Jaco as well for her warm hospitality.

I am particularly grateful to Dr. Helen Faulkner who proofread my thesis drafts and won the race, and whose friendship is dearly valued. Thanks also to David Letscher and Stephan Tillmann for many fruitful discussions, and to the American Institute of Mathematics and Oklahoma State University for supporting my research.

Thanks to Sally Miller for teaching me how to draw beautiful diagrams in xfig, to Philip Walford for encouraging my work on Conjecture 6.5.9 and to Deanna Wang for doing coffee with me when I really needed it. I am grateful to my parents for their continued support, to my grandmother Bette for the love and the chocolate and to the lovely people of SODA for their hospitality and friendship. Thanks also to Debian GNU/Linux and the many people behind it for providing an exceptional operating system that made the countless hours in front of the laptop so much less difficult.

My final thanks are to Ümit Özer for his continual love, understanding and support, especially during the most difficult periods of writing up, and for deciding to offer me Turkish coffee. Mucuk!

Contents

1	Introduction	11
1.1	Normal Surfaces	12
1.2	Layered Solid Tori	16
1.3	The Program <i>Regina</i>	25
1.3.1	History	25
1.3.2	Features	26
1.4	Structure of Material	29
2	Census of Triangulations	31
2.1	Splitting the Census Algorithm	33
2.2	Generating Face Pairings	39
2.3	Generating Gluing Permutations	43
2.3.1	Pruning Gluing Permutations	45
2.4	Processing Triangulations	50
2.4.1	Elementary Moves	50
2.4.2	0-Efficiency	54
2.4.3	Special Subcomplexes	57
2.4.4	Remaining Triangulations	62
2.5	Census Statistics	64
2.6	Face Pairings Revisited	66
2.6.1	Face Pairing Graphs	67
2.6.2	Searching for Subgraphs	69
2.6.3	Face Pairing Statistics	80
2.7	Future Directions	82
2.7.1	Eliminating Face Pairings	83
2.7.2	Improved Gluing Permutation Generation	83
3	Minimal Triangulations	85
3.1	Notation	86
3.1.1	1-Manifolds and 2-Manifolds	86
3.1.2	Seifert Fibred Spaces	86
3.1.3	Orbit Spaces	87
3.1.4	Surface Bundles	88

3.2	Common Triangulation Components	89
3.2.1	Layered Chains	90
3.2.2	The Triangular Solid Torus	92
3.2.3	Square Twisted Product Spaces	93
3.2.4	Square Untwisted Product Spaces	96
3.3	Families of Closed Triangulations	101
3.3.1	Layered Lens Spaces	101
3.3.2	Layered Loops	104
3.3.3	Layered Chain Pairs	106
3.3.4	Augmented Triangular Solid Tori	108
3.3.5	Chained Triangular Solid Tori	109
3.3.6	Plugged Triangular Solid Tori	112
3.3.7	Square Surface Bundles	115
3.3.8	Square Product Pairs	116
3.4	Closed Orientable Triangulations	118
3.4.1	One Tetrahedron	119
3.4.2	Two Tetrahedra	119
3.4.3	Three Tetrahedra	120
3.4.4	Four Tetrahedra	120
3.4.5	Five Tetrahedra	120
3.4.6	Six Tetrahedra	121
3.4.7	Summary	122
3.5	Closed Non-Orientable Triangulations	122
3.5.1	Five or Fewer Tetrahedra	123
3.5.2	Six Tetrahedra	124
3.5.3	Summary	125
3.6	Normal Surfaces	125
3.7	Future Directions	126
4	Splitting Surfaces	127
4.1	Motivation	127
4.2	Diamond Representations	130
4.3	Signatures	133
4.4	Analysing Signatures	136
4.5	Census of Signatures	146
4.5.1	Small Signatures in Detail	147
4.5.2	Signatures with 3-Manifold Triangulations	148
4.6	Future Directions	148
5	Bounding the Projective Solution Space	153
5.1	Pseudo-Convexity and Faces	154
5.2	Edge Weight Space	159
5.3	Geometric Structures	166

5.4	Mappings Between Solution Spaces	172
5.5	Constructing the Solution Space	178
6	General Position Bounds	187
6.1	Definitions	187
6.2	Intersection Pieces	190
6.3	Piece Count Calculations	196
6.4	Partial Intersections	204
6.5	Region Counting Revisited	210
6.6	Future Directions	218
6.6.1	Proving Walford's Conjecture	219
6.6.2	General Position with $k \leq d - 2$	219
6.6.3	Non-General Position	220
	Bibliography	223
A	Tables of Triangulations	227
B	Tables of Normal Surfaces	233

List of Figures

1.1	Normal arcs within a single face	13
1.2	Normal discs within a single tetrahedron	13
1.3	Normal discs meeting a face in a given type of normal arc	15
1.4	Layering a tetrahedron on a boundary edge	17
1.5	A one-triangle Möbius band	17
1.6	Layering once on a Möbius band	18
1.7	A two-triangle torus	18
1.8	A $(2, 3, 5)$ curve on a torus	19
1.9	Pushing a curve across a torus edge	19
1.10	Generators of the fundamental group of the torus	20
1.11	Slicing a torus along a non-trivial curve	20
1.12	Standard representation of a $(3, 1, -4)$ torus curve	21
1.13	The degenerate $LST(1, 1, -2)$	22
1.14	A $(1, 2, -3)$ layered solid torus	22
1.15	Adding a new tetrahedron to a layered solid torus	23
2.1	An edge of degree two belonging to only one tetrahedron	47
2.2	An edge of degree two belonging to two tetrahedra	47
2.3	An edge of degree one	48
2.4	An edge of degree one and an adjoining tetrahedron	48
2.5	A 3-2 Pachner move	50
2.6	A 2-0 vertex move	51
2.7	A 2-0 vertex move explained	51
2.8	A 2-0 edge move	52
2.9	The intermediate stages of a 2-0 edge move	52
2.10	A 2-1 edge move	53
2.11	The intermediate stage of a 2-1 edge move	54
2.12	A face with two edges identified to form a cone and a pinched cone	56
2.13	A face with three identified edges	58
2.14	A regular neighbourhood of a face with three identified edges	58
2.15	A sphere formed by identifying the boundaries of two triangles	59
2.16	Regions within the split triangulation T'	59
2.17	Pushing sphere S outside region R_1 in the neighbourhood of a vertex	60

2.18	A replacement two-tetrahedron triangular pillow	60
2.19	A case involving a vertex whose link is a multiply-punctured sphere	61
2.20	A sphere formed by creating two cones and identifying their boundaries	61
2.21	A replacement one-tetrahedron ball	61
2.22	A sample face pairing graph	67
2.23	All connected 4-valent graphs with 1, 2 and 3 vertices	69
2.24	A triple edge in a face pairing graph	70
2.25	Two tetrahedra to be joined along three faces	70
2.26	Edge identifications for the matching llr	71
2.27	Edge identifications for the matching $\kappa\kappa\alpha$	72
2.28	The vertex link of D and D' for the matching ccc	72
2.29	The partial vertex link for the matching lll	73
2.30	A one-ended chain in a face pairing graph	73
2.31	A double-ended chain in a face pairing graph	73
2.32	Identifying two faces of a tetrahedron	74
2.33	Appending a new tetrahedron to a layered solid torus	74
2.34	A broken double-ended chain in a face pairing graph	74
2.35	Two layered solid tori joined along a face	75
2.36	Allowable ways of joining two tetrahedra along two faces	75
2.37	Unallowable ways of joining two tetrahedra along two faces	76
2.38	A pinched ball formed from two tetrahedra	76
2.39	Removing the self-intersection of the disc bounded by e	77
2.40	A one-ended chain with a double handle in a face pairing graph	77
2.41	Tetrahedra corresponding to a one-ended chain with a double handle	78
2.42	A simplified diagram for a one-ended chain with a double handle	78
2.43	All diagrams for a one-ended chain with a double handle	79
2.44	Edge identifications corresponding to the different face identification diagrams . . .	79
2.45	Calculating the vertex link for diagram γ_1	80
2.46	All connected 4-valent graphs with 4 vertices	82
2.47	All connected 4-valent graphs with 5 vertices	82
3.1	A two-triangle annulus	90
3.2	A layered chain of length 1	91
3.3	A layered chain of length 2	91
3.4	Orienting the edges of a layered chain	92
3.5	A triangular solid torus	93
3.6	The three possible central surfaces of a square twisted product space	94
3.7	Constructing a square twisted product space	94
3.8	Constructing a square $T^2 \times I$ or a square $K^2 \times I$	95
3.9	A square $T^2 \times I$ and its core torus	96
3.10	A square non-orientable $K^2 \times I$ of type I and its core Klein bottle	96
3.11	A square non-orientable $K^2 \times I$ of type II and its core Klein bottle	97
3.12	A square orientable $K^2 \times I$ and its core Klein bottle	97

3.13	The three possible central surfaces of a square untwisted product space	97
3.14	Constructing a square untwisted product space	98
3.15	The boundary components of a square untwisted product space	99
3.16	Constructing a square $T^2 \times I$ of type I or II	99
3.17	Constructing a square $K^2 \times I$ of type I, II, III or IV	100
3.18	Constructing a layered lens space	102
3.19	A layered chain of length n	104
3.20	An untwisted layered loop of length n	105
3.21	A twisted layered loop of length n	105
3.22	A two-sided torus within an untwisted layered loop	105
3.23	A one-sided Klein bottle within a twisted layered loop	106
3.24	Constructing a layered chain pair	107
3.25	Annuli on the boundary of a triangular solid torus	108
3.26	Constructing an augmented triangular solid torus	109
3.27	Constructing a chained triangular solid torus of axial type	110
3.28	Constructing a chained triangular solid torus of major type	111
3.29	A positive and a negative layering upon the annulus S_i	113
3.30	Attaching a plug of major type	113
3.31	Attaching a plug of minor type	114
3.32	Edges $\alpha_1, \beta_1, \alpha_2$ and β_2 on the boundary of a square product space	115
3.33	Constructing the square torus bundle $T_{II 0,-1 -1,1}$	116
3.34	Edges $\alpha_1, \beta_1, \alpha_2$ and β_2 on the boundaries of two square product spaces	117
3.35	Constructing the square product pair $Q_{K_0, K_0 -1,0 1,1}$	118
3.36	The triangulation $L'_{3,1}$ of the lens space $L(3, 1)$	119
3.37	The minimal triangulation of the non-orientable 2-sphere bundle $S^2 \tilde{\times} S^1$	124
3.38	The minimal triangulations of $\mathbb{R}P^2 \times S^1$	124
4.1	A splitting surface within S^3/Q_8	128
4.2	Splitting tetrahedra into triangular prisms	129
4.3	Reconstructing a triangulation from a splitting surface	130
4.4	A diamond representation of a single quadrilateral	131
4.5	A diamond representation of a splitting surface within S^3/Q_8	131
4.6	The 16 allowable quadrilateral edge identifications	132
4.7	A splitting surface with no natural representation	133
4.8	The eight oriented quadrilateral edge identifications	134
4.9	Left and right arrows on a quadrilateral	134
4.10	Finding a path within a splitting surface of S^3/Q_8	135
4.11	A splitting surface with signature $(ab)(bc)(a)(c)$	135
4.12	Lifting left and right arrows away from a quadrilateral	139
4.13	A fragment of curves ρ' and ρ''	140
4.14	Quadrilaterals surrounding a hinge edge	142
4.15	A fragment of a vertex link in the underlying surface	143
4.16	The tetrahedron vertex pointed to by an arrow	144

4.17	Transformations involving parallel arrows	144
5.1	A pseudo-convex and non-pseudo-convex set in \mathbb{R}^2	154
5.2	A pseudo-convex and non-pseudo-convex set in \mathbb{R}^3	154
5.3	Normal arcs crossing a tetrahedron face	161
5.4	Normal discs in a tetrahedron	161
5.5	A triangulation of $L(5, 1)$	164
5.6	The projective edge weight solution space for the two-tetrahedron $L(5, 1)$	166
5.7	A sample 2-fan	167
5.8	Balanced and unbalanced 2-fans in \mathbb{R}^3	167
5.9	A planar cross-section of a balanced fan	167
5.10	Defining the edge weight map	173
5.11	Defining the standard map	173
5.12	Tetrahedra adjacent to edge e_i	174
5.13	A pair of balanced 2-fans intersecting within a 3-dimensional space	182
5.14	The intersection described by Equation 5.4 in the case where $d = 3$	185
6.1	Two fans in \mathbb{R}^3 that are not in general position	188
6.2	Original and final pieces for a pair of 2-fans	191

List of Tables

2.1	Statistics for the closed orientable census	65
2.2	Statistics for the closed non-orientable census	66
2.3	Precise face identifications corresponding to each transformation symbol	71
2.4	Frequency of undesirable structures within face pairing graphs	81
3.1	Properties of the individual square twisted product spaces	96
3.2	The uncategorised orientable triangulations $O_{6,1}, \dots, O_{6,6}$	122
3.3	Distinct triangulations and 3-manifolds in the closed orientable census	122
3.4	Frequencies of orientable triangulations from different families	123
3.5	Distinct triangulations and 3-manifolds in the closed non-orientable census	125
4.1	Statistics for the connected signature census	147
4.2	All connected signatures of order 1	148
4.3	All connected signatures of order 2	148
4.4	All connected signatures of order 3	149
4.5	All signatures of order ≤ 4 with enclosing 3-manifold triangulations	150
4.6	All signatures of order ≤ 8 with enclosing prime minimal 3-manifold triangulations	151
5.1	Comparisons of $U(d, k)$ with the rough bound 3^k for small d	184
6.1	Original and final piece counts for Figure 6.2	191
A.1	Closed orientable prime minimal triangulations	227
A.2	Closed non-orientable prime minimal triangulations	232
B.1	Normal surfaces in standard coordinates (1-tetrahedron orientable)	233
B.2	Normal surfaces in standard coordinates (2-tetrahedron orientable)	234
B.3	Normal surfaces in standard coordinates (3-tetrahedron orientable)	235
B.4	Normal surfaces in standard coordinates (4-tetrahedron orientable)	236
B.5	Normal surfaces in standard coordinates (5-tetrahedron orientable)	238
B.6	Normal surfaces in standard coordinates (6-tetrahedron orientable)	244
B.7	Normal surfaces in standard coordinates (2-tetrahedron non-orientable)	266
B.8	Normal surfaces in standard coordinates (3-tetrahedron non-orientable)	266
B.9	Normal surfaces in standard coordinates (6-tetrahedron non-orientable)	267
B.10	Normal surfaces in quadrilateral coordinates (1-tetrahedron orientable)	273

B.11 Normal surfaces in quadrilateral coordinates (2-tetrahedron orientable)	273
B.12 Normal surfaces in quadrilateral coordinates (3-tetrahedron orientable)	274
B.13 Normal surfaces in quadrilateral coordinates (4-tetrahedron orientable)	275
B.14 Normal surfaces in quadrilateral coordinates (5-tetrahedron orientable)	277
B.15 Normal surfaces in quadrilateral coordinates (6-tetrahedron orientable)	281
B.16 Normal surfaces in quadrilateral coordinates (2-tetrahedron non-orientable)	295
B.17 Normal surfaces in quadrilateral coordinates (3-tetrahedron non-orientable)	295
B.18 Normal surfaces in quadrilateral coordinates (6-tetrahedron non-orientable)	296

Chapter 1

Introduction

This thesis explores several aspects of the theory of normal surfaces and of minimal triangulations. Both of these fields have in recent years come to prominence in the area of computational topology. Normal surfaces lie at the heart of a large number of topological algorithms. Most of these algorithms however remain infeasibly slow for all but very small triangulations. Finding minimal triangulations of 3-manifolds is therefore critical for the efficient implementation of these algorithms.

Normal surfaces were first introduced by Kneser [21] in 1929 and were developed by Haken [8, 9] in the early 1960s for use in an algorithm for recognising the unknot. In addition Haken began the construction of an algorithm for solving the homeomorphism problem for a large class of 3-manifolds, the details of which were completed by Jaco and Oertel [16] and Hemion [10] in the 1980s and 1990s. Since then normal surfaces have become increasingly popular in the design of algorithms for solving a variety of topological problems.

Very little is known to date regarding the structure of minimal triangulations of 3-manifolds. Jaco and Rubinstein [17] develop a theory of 0-efficiency in which they describe properties of the normal surfaces within minimal triangulations. Martelli and Petronio [26] have very recently released a series of upper bounds upon the sizes of minimal triangulations of various Seifert fibred spaces. Despite these advances, full descriptions of minimal triangulations are not known even for lens spaces, the simplest closed orientable 3-manifolds.

Most of the results in which minimal triangulations are fully described are experimental. With the recent growth in computing power, it has been possible to write programs to construct censuses of 3-manifolds. The first such census was undertaken in 1989 by Hildebrand and Weeks [12] for hyperbolic 3-manifolds, and recent advances have been made by Matveev [29] and Martelli and Petronio [24] for closed 3-manifolds. Not only does a census of 3-manifolds offer insight into the structures of minimal triangulations, but it serves as a rich source of examples for exploring more general topological problems.

Computational support is of immense value when studying 3-manifolds and their triangulations, since examples that are complex enough to offer insight into a problem are generally too complex to analyse by hand. A prominent example is the program *SnapPea* [47], written by Weeks in the early 1990s, which forms an excellent tool for the study of hyperbolic 3-manifolds. Nevertheless, for the study of general 3-manifold triangulations very little software is available. Part of the

reason for this is the difficulty of implementing many of the current topological algorithms. A large portion of the work involved in this thesis has been in the development of software that implements these algorithms. The result is the program *Regina*, which is publicly available [3] and continually undergoing development.

This thesis contains a range of contributions to the theory of normal surfaces and minimal triangulations. In Chapters 2 and 3 a census of closed prime minimal 3-manifold triangulations is constructed. Both the underlying algorithm and the combinatorial structures of the resulting triangulations are examined in detail. Chapter 4 investigates splitting surfaces, which are normal surfaces within a triangulation with particularly interesting combinatorial and topological properties. In Chapters 5 and 6 aspects of the complexity of normal surfaces are studied.

Before proceeding into the main body of the thesis, some background material that is used throughout is covered in Sections 1.1 and 1.2. Following this Section 1.3 presents an overview of the program *Regina*. In Section 1.4 the structure of the remaining chapters is described in further detail.

1.1 Normal Surfaces

The theory of normal surfaces is a powerful tool in the study of 3-manifolds and in the development of algorithms for their analysis. Normal surfaces were introduced by Kneser [21] and further developed by Haken [8, 9] who used them to construct an algorithm for recognising the unknot. Haken furthermore began the construction of an algorithm for solving the homeomorphism problem for a certain large class of 3-manifolds. Difficulties with the methods of Haken were resolved by Jaco and Oertel [16] and by Hemion [10], leading to a finite time algorithm for determining whether two closed irreducible 3-manifolds are homeomorphic in the case in which one of these 3-manifolds contains an embedded two-sided incompressible surface.

Rubinstein [39, 40] has used both normal surfaces and the related class of almost normal surfaces to construct an algorithm for recognising the 3-sphere, and furthermore to construct an algorithm for solving the homeomorphism problem for all closed orientable irreducible 3-manifolds with Heegaard genus ≤ 2 . Jaco and Tollefson [19] have employed normal surfaces in algorithms for a variety of tasks. These include the decomposition of a closed 3-manifold into irreducible 3-manifolds and the decomposition of an orientable 3-manifold into its Jaco-Shalen-Johannson decomposition, as well as the recognition of orientable Seifert fibred spaces and handlebodies. These algorithms and others are further developed by Jaco, Letscher and Rubinstein in [15].

The simplification of 3-manifold triangulations is a field in which normal surfaces have recently enjoyed success. Jaco and Rubinstein [17, 18] develop a theory of 0-efficiency and 1-efficiency in which they employ normal spheres and tori to modify a 3-manifold triangulation so that it uses fewer tetrahedra.

We present here an outline of the core results of normal surface theory. Normal surfaces are used as a tool in Chapter 2, and in Chapters 4, 5 and 6 our attention is explicitly devoted to the study of normal surfaces and related problems. In Appendix B we present a list of the vertex normal surfaces of every closed prime minimal triangulation formed from at most six tetrahedra. For a more extensive review of normal surface theory, the reader is referred to [10].

Definition 1.1.1 (Normal Arc) Let F be a face in a 3-manifold triangulation. A *normal arc* on F is an embedded arc whose interior lies in the interior of F and whose endpoints lie on distinct edges of F . A normal arc may not meet any of the vertices of the face.

Each face then allows for three types of normal arc, corresponding to the three choices of edges on which the endpoints may lie. Examples of each of these three types of normal arc are illustrated in Figure 1.1.




Figure 1.1: Normal arcs within a single face

Definition 1.1.2 (Normal Disc) Let Δ be a tetrahedron in a 3-manifold triangulation. A *normal disc* in Δ is an embedded disc whose interior lies in the interior of Δ and whose boundary consists of a sequence of normal arcs contained in distinct faces of Δ .

Each tetrahedron allows for seven types of normal disc. Four of these disc types correspond to triangular discs surrounding each of the four vertices. The remaining three disc types correspond to quadrilateral discs, each of which passes through all four faces of Δ and separates two opposite edges of Δ . The four types of triangular disc and one of the three types of quadrilateral disc are illustrated in Figure 1.2.




Figure 1.2: Normal discs within a single tetrahedron

Definition 1.1.3 (Embedded Normal Surface) Let T be a 3-manifold triangulation. An *embedded normal surface* within T is a properly embedded surface within T meeting each tetrahedron in a collection of disjoint normal discs. All normal surfaces within this thesis are assumed to be embedded unless otherwise indicated.

It is known that for many properties P in which we are interested, the existence of an embedded surface with property P within a 3-manifold triangulation implies the existence of an embedded normal surface with this same property P . Examples of such existence theorems due to Haken [8], Kneser [21] and Schubert [42] are as follows.

Theorem 1.1.4 If a 3-manifold triangulation contains a properly embedded essential disc, then it contains an essential embedded normal disc.

Theorem 1.1.5 If a 3-manifold triangulation contains an embedded essential 2-sphere, then it contains an essential embedded normal 2-sphere.

By restricting our attention from the embedded surfaces within a triangulation to the embedded normal surfaces within a triangulation, we can therefore take advantage of the additional structure imposed upon each surface. In particular, we can convert many topological problems into algebraic problems as follows.

Definition 1.1.6 (Vector Representation) Consider a 3-manifold triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$. Any normal surface S within this triangulation can be represented as a vector in \mathbb{R}^{7n} . Specifically, we define the *vector representation* of S to be the vector

$$\mathbf{v} = (t_{1,1}, t_{1,2}, t_{1,3}, t_{1,4}, q_{1,1}, q_{1,2}, q_{1,3} ; t_{2,1}, t_{2,2}, t_{2,3}, t_{2,4}, q_{2,1}, q_{2,2}, q_{2,3} ; \dots, q_{n,3}), \quad (1.1)$$

where the individual coordinates of \mathbf{v} are defined as follows.

- For each tetrahedron Δ_k , the coordinates $t_{k,1}, \dots, t_{k,4}$ correspond to the four types of triangular normal disc in Δ_k . Specifically, $t_{k,i}$ represents the number of triangular discs in Δ_k of the i th type.
- Similarly, for each tetrahedron Δ_k the coordinates $q_{k,1}, \dots, q_{k,3}$ correspond to the three types of quadrilateral normal disc in Δ_k . The individual coordinate $q_{k,i}$ represents the number of quadrilateral discs in Δ_k of the i th type.

By analysing the structure of a normal surface within a triangulation, we can derive constraints upon its vector representation. For instance, any two quadrilaterals of different types within a tetrahedron must intersect. Thus if S is an embedded normal surface then for each tetrahedron Δ_k there is at most one non-zero quadrilateral coordinate $q_{k,i}$.

We can derive an entire family of linear equations by analysing normal arcs on the faces of the triangulation as follows.

Theorem 1.1.7 (Matching Equations) Consider a triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$. Let S be an embedded normal surface within this triangulation with vector representation \mathbf{v} as described in Equation 1.1. Let F be a non-boundary face in this triangulation and consider any one of the three types of normal arc on F . Denote this arc type by f .

Let Δ_k and $\Delta_{k'}$ be the tetrahedra on either side of face F . Note that precisely one of the four triangular disc types and precisely one of the three quadrilateral disc types within Δ_k meet F in a normal arc of type f . Let $t_{k,r}$ and $q_{k,s}$ be the corresponding coordinates of \mathbf{v} . Similarly, let $t_{k',r'}$ and $q_{k',s'}$ be the coordinates of \mathbf{v} corresponding to the normal disc types within $\Delta_{k'}$ that meet F in a normal arc of type f . Then it is true that $t_{k,r} + q_{k,s} = t_{k',r'} + q_{k',s'}$.

Proof By counting normal arcs on the faces of Δ_k we see that face F must contain precisely $t_{k,r} + q_{k,s}$ normal arcs of type f . Similarly we can count normal arcs on the faces of $\Delta_{k'}$ to conclude that face F contains precisely $t_{k',r'} + q_{k',s'}$ normal arcs of type f . ■

Figure 1.3 presents an illustration of this scenario in which Δ_k has corresponding coordinates $t_{k,r} = q_{k,s} = 1$ and $\Delta_{k'}$ has corresponding coordinates $t_{k',r'} = 2$ and $q_{k',s'} = 0$. Note that all four normal discs meet the interior face in the same type of normal arc.




Figure 1.3: Normal discs meeting a face in a given type of normal arc

The equivalence between the topological problem of analysing normal surfaces and the algebraic problem of analysing their vector representations is shown in the following result of Haken [8, 9].

Theorem 1.1.8 Consider once more a triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$. Let \mathbf{v} be a vector in \mathbb{R}^{7n} as described by Equation 1.1 for which the following conditions are satisfied.

- \mathbf{v} is not the zero vector;
- Each coordinate of \mathbf{v} is a non-negative integer;
- The vector \mathbf{v} satisfies the matching equations as described in Theorem 1.1.7;
- For each tetrahedron Δ_k at most one of the quadrilateral coordinates $q_{k,1}$, $q_{k,2}$ and $q_{k,3}$ is non-zero.

Then there is an embedded normal surface (or possibly a union of disjoint embedded normal surfaces) within the triangulation whose vector representation is \mathbf{v} . Furthermore, this surface or union of surfaces is unique up to normal isotopy, i.e., an isotopy that is invariant upon each tetrahedron Δ_k .

Normal surfaces are frequently used within algorithms as a means for testing whether a particular type of embedded surface exists. There are however generally infinitely many vectors of \mathbb{R}^{7n} that satisfy the conditions of Theorem 1.1.8. We thus describe two restricted classes of normal surfaces that can be used as bases for all normal surfaces within a triangulation.

Definitions 1.1.9 (Fundamental and Vertex Surfaces) Consider the normal surfaces within some 3-manifold triangulation. If S_1 and S_2 are surfaces with vector representations \mathbf{v}_1 and \mathbf{v}_2 respectively then we let $S_1 + S_2$ denote the surface whose vector representation is $\mathbf{v}_1 + \mathbf{v}_2$.

Surface S is called a *fundamental surface* if it cannot be written as the sum $S = S_1 + S_2$ for any non-empty surfaces S_1 and S_2 . Surface S is called a *vertex surface* if the equation $kS = lS_1 + mS_2$ cannot be satisfied for any positive integers k , l and m and any non-empty surfaces S_1 and S_2 that are not simply multiples of S .

It is clear that for any positive integer k the surface kS is a vertex surface if and only if S itself is a vertex surface. We thus define the *class of surfaces* $[S]$ to be the set of all surfaces that are rational multiples of the surface S .

An important observation is that there are only finitely many classes of vertex embedded normal surfaces for any 3-manifold triangulation, and that using the linear equations and inequalities described in Theorem 1.1.8 these classes can be enumerated using the techniques of linear programming.

In order that each class of surfaces can be assigned a unique representative vector, we define the projective solution space as follows.

Definition 1.1.10 (Projective Solution Space) Let T be a 3-manifold triangulation formed from n tetrahedra. The *projective solution space* of T is defined to be the set of all vectors $\mathbf{v} \in \mathbb{R}^{7n}$ for which the following constraints are satisfied.

- \mathbf{v} satisfies the matching equations described in Theorem 1.1.7;
- The coordinates of \mathbf{v} are each non-negative;
- The coordinates of \mathbf{v} sum to 1.

It is straightforward to show that the projective solution space forms a finite convex polytope embedded in \mathbb{R}^{7n} . Furthermore, the classes of vertex embedded normal surfaces correspond precisely to the vertices of this polytope for which at most one quadrilateral coordinate for each tetrahedron is non-zero. Note that the coordinates of each vertex of this polytope lie in the range $[0, 1]$, and in particular that these vertices are generally not integer vectors.

We can thus use normal surfaces in a variety of algorithms as follows. Suppose that we seek to determine whether a 3-manifold triangulation T contains an embedded surface satisfying some property P . We begin by proving that if such a surface exists then there is an embedded normal surface with property P . We then construct an algorithm for determining from the vertex embedded normal surfaces of T whether such an embedded normal surface exists. An example of such an algorithm is presented in Algorithm 2.4.13 in which we use vertex normal surfaces to determine the existence or otherwise of an embedded normal 2-sphere. The vertex normal surfaces themselves are finally enumerated using the linear programming techniques mentioned above.

Note that the enumeration of vertex normal surfaces is very slow, and that the running time of all current algorithms is at least exponential in the number of tetrahedra. In [45] Tollefson presents an alternative coordinate system using only the quadrilateral coordinates $\{q_{k,i}\}$, under which normal surfaces are represented as vectors in \mathbb{R}^{3n} instead of vectors in \mathbb{R}^{7n} . This reduction in the dimension of the underlying vector space can lead to dramatic improvements in the efficiency of normal surface algorithms.

1.2 Layered Solid Tori

In this section we describe in detail the construction of a layered solid torus, a building block that appears frequently within 3-manifold triangulations. Layered solid tori appear in some of the proofs of Chapter 2 and are used heavily in the constructions described in Chapter 3.

Layered solid tori have been discussed in a variety of informal contexts by Jaco and Rubinstein and are described by these authors in [17]. Analogous constructs involving special spines of 3-manifolds have been described in detail by Matveev [29] and by Martelli and Petronio [26].

The results in this section are public knowledge. Some such as the torus curve results are commonplace in the literature. Others cannot necessarily be found written formally but have been discussed frequently by Jaco, Letscher and Rubinstein in informal contexts. Attention must be drawn however to Theorem 1.2.19 with which this section concludes. Whilst the result is certainly not new, the proof is original and illustrates a different approach to a standard problem.

We begin then by describing the process of layering. Layering is a transformation that, when applied to a triangulation with boundary, does not change the underlying 3-manifold but does change the curves formed by the boundary edges of the triangulation.

Definition 1.2.1 (Layering) Let M be a triangulation with boundary and let e be one of its boundary edges. To *layer a tetrahedron on edge e* , or just to *layer on edge e* , is to take a new tetrahedron T , choose two of its faces and identify them with the two boundary faces on either side of e without twists. This procedure is illustrated in Figure 1.4.




Figure 1.4: Layering a tetrahedron on a boundary edge

Note that layering on a boundary edge does not change the underlying 3-manifold; the only effect is to thicken the boundary around edge e . Note however that edge e is no longer a boundary edge, and instead edge f (which in general represents a different curve on the underlying 3-manifold boundary) has been added as a new boundary edge.

Definition 1.2.2 (Layered Solid Torus) A *layered solid torus* is a triangulation of a solid torus formed as follows. We begin with the Möbius band illustrated in Figure 1.5, where the two edges marked e are identified according to the arrows and where edge g is a boundary edge. If we thicken this Möbius band slightly, we can imagine it as a solid torus with two boundary faces, one on each side of the Möbius band. We then begin layering upon boundary edges one at a time. We may layer as many times as we like or we may layer no times at all.




Figure 1.5: A one-triangle Möbius band

The two figures produced by layering on edges e and g (separately) are illustrated in Figure 1.6. The first of these diagrams, which represents a layering on edge e , depicts the new tetrahedron cut in half. This tetrahedron in fact sits on top of the Möbius strip, runs off the diagram to the right and returns from the left to simultaneously sit beneath the Möbius strip. The second diagram, which represents a layering on edge g , has the new tetrahedron sitting on top of the Möbius strip, running around and below edge g and returning to simultaneously sit beneath the Möbius strip.




Figure 1.6: Layering once on a Möbius band

Once we have layered on edge e , it is no longer necessary to thicken the Möbius band; instead the figure forms a triangulation in its own right (still of a solid torus) with the Möbius band reduced to an internal face of this triangulation.

Until that point however, we do not yet have a proper triangulation since some boundary edges are still pinched together (as in the second diagram of Figure 1.6). Such a triangulation is called a *degenerate layered solid torus* and is still of interest, since although it cannot be a triangulation on its own, it may still occur as a subcomplex of some larger triangulation.

We can observe that each layered solid torus has two boundary faces and represents the same underlying 3-manifold, i.e., a solid torus. What distinguishes the different layered solid tori, and in particular what ensures that the larger layered solid tori remain useful, is the different patterns of curves that their boundary edges make upon the boundary torus.

Definition 1.2.3 (Three-Parameter Torus Curves) Let T be a torus formed from two triangles, as illustrated in Figure 1.7. Label one of these triangles $+$ and the other $-$. Select some ordering of the three edges and label these edges e_1 , e_2 and e_3 accordingly.




Figure 1.7: A two-triangle torus

Consider some oriented closed curve on this torus in general position with respect to each edge (i.e., not tangential to any edge), as illustrated in Figure 1.8. From this curve we can assign a number to each edge of the torus, this being the number of times the curve crosses the edge from $+$ to $-$ minus the number of times it crosses the edge from $-$ to $+$.




Figure 1.8: A $(2, 3, 5)$ curve on a torus

If the numbers assigned to edges e_1 , e_2 and e_3 are p , q and r respectively, we refer to our oriented curve as a (p, q, r) curve. Thus, for instance, the curve in Figure 1.8 is a $(2, 3, -5)$ curve.

Lemma 1.2.4 For any (p, q, r) curve on a torus, $p + q + r = 0$.

Proof Since the three torus edges form the boundary of face $+$, it follows that $p + q + r$ is the total number of times the curve leaves face $+$ minus the total number of times it enters face $+$. Since the curve is closed, this total must be zero. ■

Lemma 1.2.5 (Uniqueness of Curve Parameters) If a (p, q, r) curve on a torus is homotopy equivalent to a (p', q', r') curve, then $p = p'$, $q = q'$ and $r = r'$.

Proof This is seen from the following facts. Each edge of the torus is in fact a loop, and so a curve cannot be pushed off the end of an edge (thereby reducing the number of intersections with that edge). Furthermore, pushing a segment of the curve across an edge as illustrated in Figure 1.9 does not alter the parameters (p, q, r) , since the two new intersections add $+1$ and -1 to the count for edge e , resulting in no overall change.




Figure 1.9: Pushing a curve across a torus edge

Lemma 1.2.6 Select elements α and β from the fundamental group of the torus. Let α be a (p, q, r) curve on the torus and let β be a (p', q', r') curve, noting from Lemma 1.2.5 that these parameters are well-defined. Then element $\alpha \cdot \beta$ in the fundamental group is a $(p + p', q + q', r + r')$ curve on the torus.

Proof This follows immediately from the observation that the intersections of $\alpha \cdot \beta$ with a torus edge are precisely the intersections of α with the edge combined with the intersections of β with the edge. ■

Lemma 1.2.7 (Uniqueness of Torus Curves) For any integer triple (p, q, r) with $p + q + r = 0$, there exist (p, q, r) torus curves on a torus. Furthermore, all such curves are homotopy equivalent.

Proof In Figure 1.10, closed curves α and β are drawn; the fundamental group of the torus is known to be the free abelian group generated by these two curves.




Figure 1.10: Generators of the fundamental group of the torus

It can be seen from the diagram that α and β are $(1, 0, -1)$ and $(0, 1, -1)$ curves on the torus. Hence, from Lemma 1.2.6 and the fact that $r = -p - q$, we see that the curve $\alpha^p \cdot \beta^q$ is a (p, q, r) torus curve. Thus such curves exist.

Furthermore, any (p, q, r) torus curve must be homotopy equivalent to $\alpha^x \cdot \beta^y$ for some x, y , and is thus homotopy equivalent to an $(x, y, -x - y)$ torus curve. Therefore $x = p$ and $y = q$, and so we see that all (p, q, r) curves are homotopy equivalent. ■

Lemma 1.2.8 (Embedded Torus Curves) If a (p, q, r) curve on a torus does not intersect itself then either p, q and r are pairwise coprime or $(p, q, r) = (0, 0, 0)$.

Proof Consider some (p, q, r) curve κ that does not intersect itself. If $(p, q, r) \neq (0, 0, 0)$ then Lemma 1.2.5 implies that κ cannot bound a disc on the torus, and so if we slice our torus along κ we obtain an annulus as illustrated in Figure 1.11.




Figure 1.11: Slicing a torus along a non-trivial curve

It follows then that κ and λ as illustrated in the figure together generate the fundamental group of the torus, and in particular that κ cannot be expressed as μ^d in the fundamental group for any other curve μ and any $d > 1$.

Suppose then that p, q and r are not pairwise coprime. Without loss of generality let $\gcd(p, q) = d > 1$. From Lemma 1.2.4 we see that $\gcd(p, q, r) = d$ and so $(p, q, r) = d(p', q', r')$ for some integers p', q' and r' . Using Lemma 1.2.7 we see that some (p', q', r') curve exists. Call this curve μ ; then Lemma 1.2.6 shows that $\kappa = \mu^d$, contradicting our earlier conclusion.

Thanks must go to Rubinstein for suggesting this particularly nice approach. ■

Theorem 1.2.9 (Standard Representation of Curves) Let p, q and r be any pairwise coprime integers for which $p + q + r = 0$.

Consider the \mathbb{R}^2 cover of the torus, in which two points (x, y) and (x', y') in \mathbb{R}^2 map to the same point of the torus precisely when both $x - x'$ and $y - y'$ are integers. In particular, each

unit square in this cover represents a single copy of the torus, with the top edge identified with the bottom and the left edge identified with the right in the usual way.

Consider the line $qx + py = 0$ in this \mathbb{R}^2 cover with an orientation pointing in the direction from $(0,0)$ to $(-p,q)$. The image of this line on the torus is an oriented closed curve with no self-intersections, and is in fact homotopy equivalent to a (p,q,r) curve. Furthermore, the segment of this line between $(0,0)$ and $(-p,q)$ (including the first endpoint but excluding the second) is in 1-1 correspondence with the image of the line on the torus.

We call this representation of a (p,q,r) curve on the torus the *standard representation* of a (p,q,r) curve.

Proof Call the line on the \mathbb{R}^2 cover l . Since p and q are coprime, l contains no integer points between $(0,0)$ and $(-p,q)$. Thus two points on l map to the same point of the torus if and only if they are separated by some multiple of $(-p,q)$, and hence the segment from $(0,0)$ to $(-p,q)$ is in 1-1 correspondence with the image of l as claimed.

Since $(0,0)$ and $(-p,q)$ map to the same point, this image is indeed a closed curve. Since the gradient of the line is constant, this image has no self intersections.

Finally, according to Figure 1.10, the segment from $(0,0)$ to $(-p,q)$ is homotopy equivalent to p copies of α followed by q copies of β . From Lemma 1.2.6 then, we see that we in fact have a $p(1,0,-1) + q(0,1,-1) = (p,q,-p-q) = (p,q,r)$ curve on the torus. ■

Example 1.2.10 The construction of the standard representation of a $(3,1,-4)$ torus curve is seen in Figure 1.12. Note that the fourth intersection with edge e_3 is difficult to see; it is in fact where the curve passes through the torus vertex.




Figure 1.12: Standard representation of a $(3,1,-4)$ torus curve

Definition 1.2.11 (Layered Solid Torus Parameters) Let L be a layered solid torus. On the two faces that form the boundary torus, draw the boundary of a meridinal disc of the underlying solid torus. Assign this meridinal curve some arbitrary orientation and arbitrarily label the two boundary faces $+$ and $-$.

The meridinal curve then forms some (p,q,r) curve on the boundary torus. p , q and r are said to be the *parameters* of layered solid torus L , and L is said to be a (p,q,r) *layered solid torus*, denoted $LST(p,q,r)$.

Lemma 1.2.12 A (p,q,r) layered solid torus is also an (x,y,z) layered solid torus if and only if the triple (p,q,r) can be rearranged to form either (x,y,z) or $(-x,-y,-z)$.

Proof The only ways in which we can calculate the layered solid torus parameters in a different fashion are to select the boundary edges in a different order, to reverse the labels $+$ and $-$, and to reverse the orientation of the meridinal curve. These changes lead precisely to the equivalences described above. ■

Example 1.2.13 The degenerate layered solid torus with zero tetrahedra, i.e., the Möbius band illustrated in Figure 1.13, is a $(1, 1, -2)$ layered solid torus. The meridinal curve is illustrated by the dotted line. The three edges of the (degenerate) boundary torus that it meets are g and the two sides of the identified edges marked e ; these intersections are marked by black circles in the diagram.




Figure 1.13: The degenerate LST(1, 1, -2)

Example 1.2.14 A $(1, 2, -3)$ layered solid torus made from one tetrahedron is illustrated in Figure 1.14; note that this is formed by layering upon edge e from Figure 1.13. The back two faces of the tetrahedron are identified; specifically face PQR is identified with face QRS . The meridinal curve is again illustrated by the dotted line drawn on the boundary faces (the front two faces) and its intersections with the boundary torus edges are marked.




Figure 1.14: A $(1, 2, -3)$ layered solid torus

Lemma 1.2.15 In any (p, q, r) layered solid torus, the integers p , q and r are pairwise coprime.

Proof Since any solid torus has an embedded meridinal disc, a meridinal curve with no self intersections can be selected. The trivial $(0, 0, 0)$ curve can never be a meridinal curve and so the result follows from Lemma 1.2.8. ■

Lemma 1.2.16 Layering on the edge with parameter r in a (p, q, r) layered solid torus produces a $(p, -q, q - p)$ layered solid torus.




Figure 1.15: Adding a new tetrahedron to a layered solid torus

Proof The layering described is illustrated in Figure 1.15, where edges e_1 , e_2 and e_3 have parameters p , q and r respectively and the boundary faces of the original layered solid torus have been marked $+$ and $-$.

When the new tetrahedron is added, two new boundary faces are produced. These have also been marked $+$ and $-$, and are bordered by edges e_1 , e_2 and e_4 . The important fact here is that edges e_1 and e_2 belong to both the old torus boundary and the new torus boundary.

Thus the number of times the meridinal disc crosses e_1 from $+$ to $-$ in the old torus boundary is the same as the number of times it crosses e_1 from $+$ to $-$ in the new torus boundary. Similarly, the number of times the meridinal disc crosses e_2 from $+$ to $-$ in the old torus boundary is in fact the number of times it crosses e_2 from $-$ to $+$ in the new torus boundary (since in the new boundary, faces $+$ and $-$ have moved to opposite sides of edge e_2).

In the new layered solid torus therefore, edges e_1 and e_2 have parameters p and $-q$ respectively. Since the three edge parameters add to zero by Lemma 1.2.4, it follows that the new layered solid torus is an $\text{LST}(p, -q, q - p)$. ■

Algorithm 1.2.17 (Layered Solid Torus Construction) For any pairwise coprime integers p , q and r satisfying $p + q + r = 0$, the following algorithm can be used to construct a layered solid torus with parameters (p, q, r) . Without loss of generality, let $|r| \geq |p|, |q|$.

Note that if $|r| > 1$ then we in fact have $|r| > |p|, |q|$, since otherwise we would have $p = \pm r$ or $q = \pm r$, contradicting Lemma 1.2.15 which requires the three parameters to be pairwise coprime.

Suppose that $|r| \geq 3$. In this case, since $|r| > |p|, |q|$, equation $p + q + r = 0$ tells us that r has the opposite sign to both p and q and that $|p| + |q| = |r|$. We use this algorithm to construct an $\text{LST}(p, -q, q - p)$ and then layer on the edge with parameter $|q - p|$ to obtain an $\text{LST}(p, q, r)$, as seen in Lemma 1.2.16.

Since $\max(|p|, |q|, |r|)$ decreases at each step, we are eventually reduced to a case in which $|r| \leq 2$. The only pairwise coprime triples that satisfy this condition are $(1, 1, -2)$ and $(0, 1, -1)$. To construct an $\text{LST}(1, 1, -2)$ we simply use a Möbius band (see Example 1.2.13), and to construct an $\text{LST}(0, 1, -1)$ we start with the Möbius band and layer on the edge with parameter -2 .

Recall from Definition 1.2.2 that the only degenerate layered solid tori are those in which we have not yet layered upon one of the original Möbius band edges with parameter ± 1 . The only cases therefore in which this algorithm produces a degenerate layered solid torus are for the $\text{LST}(1, 1, -2)$ and the $\text{LST}(0, 1, -1)$.

To construct a non-degenerate $\text{LST}(1, 1, -2)$, we start with a Möbius band, i.e., a degenerate $\text{LST}(1, 1, -2)$, layer on one of the parameter 1 edges to produce a non-degenerate $\text{LST}(1, 2, -3)$ and then layer on the edge with parameter -3 to produce a non-degenerate $\text{LST}(1, 1, -2)$.

To construct a non-degenerate $\text{LST}(0, 1, -1)$, we start with a Möbius band, layer on the edge with parameter -2 to produce a degenerate $\text{LST}(0, 1, -1)$ and then layer on one of the edges with parameter ± 1 to produce a non-degenerate $\text{LST}(0, 1, -1)$.

Example 1.2.18 In this example we construct an $\text{LST}(3, 7, -10)$ using Algorithm 1.2.17.

To form an $\text{LST}(3, 7, -10)$ we layer on edge 4 in an $\text{LST}(3, 4, -7)$. To form an $\text{LST}(3, 4, -7)$ we layer on edge 1 in an $\text{LST}(1, 3, -4)$. To form an $\text{LST}(1, 3, -4)$ we layer on edge 2 in an $\text{LST}(1, 2, -3)$, and to form an $\text{LST}(1, 2, -3)$ we layer on one of the parameter 1 edges in the Möbius band $\text{LST}(1, 1, -2)$. The resulting triangulation has four tetrahedra.

Since we use layered solid tori as building blocks of 3-manifold triangulations, it is important to understand how curves on a layered solid torus match up with curves on the adjoining pieces of a 3-manifold, and to be able to identify other important curves on the torus boundary (such as longitudes).

Theorem 1.2.19 (Intersections of Torus Curves) Let (p, q, r) be pairwise coprime integers with $p + q + r = 0$ and let (p', q', r') be pairwise coprime integers with $p' + q' + r' = 0$. Let K be the number of times the standard representations of a (p, q, r) curve and a (p', q', r') curve on a torus intersect. Then

$$K = \left| \det \begin{pmatrix} p & q \\ p' & q' \end{pmatrix} \right| = \left| \det \begin{pmatrix} q & r \\ q' & r' \end{pmatrix} \right| = \left| \det \begin{pmatrix} r & p \\ r' & p' \end{pmatrix} \right|. \quad (1.2)$$

Proof Consider the \mathbb{R}^2 cover of the torus as described in Theorem 1.2.9. From this theorem, recall that the standard representations of our two torus curves are the lines $qx + py = 0$ and $q'x + p'y = 0$ on this \mathbb{R}^2 cover.

Recall also that a unit square in the cover represents a single copy of the torus. Thus the unit square $[0, 1) \times [0, 1) \subset \mathbb{R}^2$ maps to each point of the torus once and only once.

It follows that a point $(x, y) \in [0, 1) \times [0, 1)$ represents a point of intersection of our two torus curves if and only if it is equivalent to a point on the line $qx + py = 0$ and to a point on the line $q'x + p'y = 0$. This happens if and only if $qx + py \in \mathbb{Z}$ and $q'x + p'y \in \mathbb{Z}$.

Thus the number of intersections K is the number of points $(x, y) \in [0, 1) \times [0, 1)$ for which $(qx + py, q'x + p'y) \in \mathbb{Z} \times \mathbb{Z}$, i.e.,

$$\begin{bmatrix} q & p \\ q' & p' \end{bmatrix} \times \begin{bmatrix} x \\ y \end{bmatrix} \in \mathbb{Z} \times \mathbb{Z}. \quad (1.3)$$

Since each point on an edge of our unit square is equivalent to a point on the opposite edge and

all four vertices of the unit square are equivalent, we can instead consider the closed unit square $[0, 1] \times [0, 1]$ and observe that each solution on the boundary is being counted twice except for the vertex $(0, 0)$ which is being counted four times. Thus $K = I + \frac{1}{2}B - 1$, where I and B are the numbers of points in the square interior and on the square boundary that satisfy Equation 1.3 (the -1 taking care of the vertex which is still being counted twice in $\frac{1}{2}B$ and which is always a solution).

Consider now the parallelogram formed by applying the affine transformation

$$\begin{bmatrix} q & p \\ q' & p' \end{bmatrix}$$

to the unit square $[0, 1] \times [0, 1]$. It follows then that I and B are the numbers of lattice points in the interior and on the boundary of this parallelogram respectively.

By Pick's Theorem [35] however, the area of a polygon whose vertices are all lattice points is precisely $I + \frac{1}{2}B - 1$, where I and B are the numbers of lattice points in the interior and on the boundary of the polygon. Thus K is simply the area of this parallelogram, which is

$$\left| \det \begin{pmatrix} q & p \\ q' & p' \end{pmatrix} \right|$$

since the area of the unit square is 1.

Finally, since $r = p + q$ and $r' = p' + q'$, it is simple to see that the three determinants in Equation 1.2 are all equal. ■

1.3 The Program *Regina*

Much of the work presented in this thesis has required an extensive level of computational support. To this end, the program *Regina* [3] has been written and heavily utilised.

Regina is publicly available under the GNU General Public License and can be downloaded from <http://regina.sourceforge.net/>. Both the user interface and the underlying programmer interface are thoroughly documented.

1.3.1 History

The program that has become *Regina* began as a joint project of David Letscher, Richard Rannard and the author in 1997 under the guidance of J. Hyam Rubinstein, with the aim of implementing a variety of normal surface algorithms. Much planning was done but very little was written.

In early 1999 Letscher revived the project on his own and presented *Normal*, a program written in Java that would simplify triangulations and find vertex normal surfaces. This program however was intended as a proof-of-concept only, and later that year Letscher and this author began writing a new program from scratch. The new program was written in C++ and carefully designed for rigour and extensibility.

Although Letscher did not have time to continue with the project, he continues to offer much appreciated technical advice. This author took over the maintenance of the project in early 2000,

and in late 2000 the program, now called *Regina*, saw its first public release.

Regina continues to grow and currently enjoys a new release every few months. Today almost all of the *Regina* source code has been written by this author, though not without invaluable advice from many others. In particular, Marc Culler, Nathan Dunfield, William Jaco, David Letscher, J. Hyam Rubinstein and Jeff Weeks should be thanked for many fruitful discussions.

1.3.2 Features

At the time of writing, *Regina* contains approximately 41,000 physical source lines of code (i.e., lines of code that are neither blank nor comment) and approximately 38,000 additional lines of comment within this source code (still excluding the users' reference manual).¹ A list of the more noteworthy features of *Regina* is presented below.

Triangulations

The fundamental objects with which a user works when running *Regina* are 3-manifold triangulations. The program supports the following methods of creating triangulations.

- Automated generation of standard triangulations such as layered solid tori (described in Section 1.2) and layered lens spaces (described in Section 3.3.1);
- Importing triangulations saved from *SnapPea* [47] (the hyperbolic 3-manifold software written by Jeffrey Weeks) and parsing dehydration strings (text-based representations of triangulations defined and used by Callahan, Hildebrand and Weeks in [5]);
- Manually building triangulations by entering the individual identifications between tetrahedron faces by hand.

Once a triangulation has been created or imported, a user will generally wish to analyse it and perhaps modify it. For the analysis of a triangulation, the properties that can be computed by the software include the following.

- Detailed combinatorial information about the triangulation, its skeleton and its boundary components, including vertex links and the shapes formed by the various triangulation faces;
- A variety of homology and homotopy groups;
- Attributes related to the normal surfaces within a triangulation such as 0-efficiency (described in Section 2.4.2) and the existence of splitting surfaces (described in Section 4.1).

In addition the software contains a variety of recognition routines for detecting particular structures within a triangulation. These routines recognise common building blocks such as layered solid tori (described in Section 1.2) and the various components discussed in Section 3.2, and also detect members of most of the infinite families described in Section 3.3. In particular this means that *Regina* can frequently recognise the underlying 3-manifolds for well-structured triangulations that it has not previously encountered, simply from their combinatorial structures.

For the modification of a triangulation, the following procedures are available.

¹Estimates were obtained using the software metrics tools *SLOCCount* by David A. Wheeler (<http://www.dwheeler.com/sloccount/>) and *CCCC* by Tim Littlefair (<http://cccc.sourceforge.net/>).

- Elementary moves including the operations described in Section 2.4.1 and a variety of similar transformations that are local to a small number of tetrahedra, many of which were suggested by Letscher;
- Automated simplification in which the program attempts to use a combination of these elementary moves to reduce the number of tetrahedra in a triangulation as far as possible, though there is no guarantee that the smallest possible number of tetrahedra will be achieved;
- Barycentric subdivision and the truncation of ideal vertices (vertices whose links are neither 2-spheres nor discs);
- Conversion of a non-orientable triangulation to an orientable double cover.

Census Creation

Regina can be used to form a census of all 3-manifold triangulations satisfying a particular set of census constraints. Chapter 2 is devoted to describing the details of the underlying algorithm and Chapter 3 discusses the results of two such censuses, though the program supports a wider variety of census constraints than those described in Chapters 2 and 3.

Normal Surfaces

Normal surfaces are outlined in Section 1.1, and were the original motivation for writing this software. The program is therefore capable of enumerating all vertex normal surfaces or almost normal surfaces² of a triangulation, an operation required by most high-level topological algorithms that rely upon normal surface theory.

This vertex enumeration can be performed in a variety of coordinate systems. For an n -tetrahedron triangulation this includes the $7n$ standard triangle and quadrilateral coordinates described in Section 1.1 as well as the smaller set of $3n$ quadrilateral-only coordinates introduced by Tollefson [45] for algorithmic efficiency. The enumeration can be restricted to embedded normal surfaces or can be expanded to include immersed and singular surfaces. Furthermore elementary support is present for spun normal surfaces, discussed in detail by Tillmann [44], which are non-compact surfaces with infinitely many discs found in ideal triangulations (triangulations whose vertex links have non-trivial genus).

For the analysis of normal surfaces, *Regina* offers the following tools.

- Viewing normal surfaces in a variety of coordinate systems, including the standard and quadrilateral-only coordinates discussed above as well as the edge weight coordinates described in Section 5.2;
- Calculating basic properties of normal surfaces such as Euler characteristic, orientability and one-sidedness;
- Recognising standard surfaces within a triangulation such as splitting surfaces (described in Section 4.1) and vertex and edge links;

²Almost normal surfaces are closely related to normal surfaces and are used by Rubinstein in his 3-sphere recognition algorithm [39, 40].

- Filtering large lists of normal surfaces by various properties such as Euler characteristic and orientability.

In addition the program can crush a normal surface to a point within a triangulation. Crushing is a powerful tool for the analysis of the role played by a surface within a 3-manifold. This transformation is used within Jaco and Rubinstein’s 0-efficiency algorithm as presented in [17], and outlined in the proof of Theorem 2.4.6.

A related but significantly more complex procedure is the cutting open of a triangulation along a normal surface and the retriangulation of the resulting 3-manifold(s), introduced by Haken [9] for solving the homeomorphism problem for 3-manifolds containing incompressible surfaces and used in a variety of algorithms since. Like crushing, this procedure is a powerful tool and furthermore avoids the difficulties suffered by the crushing process in which topological information can be lost or (in the non-orientable case) invalid 3-manifold triangulations can be created.

Cutting along a surface however is more difficult than crushing since the individual tetrahedra containing the normal discs can become heavily subdivided. The many potential combinations of discs lead to many different ways in which this subdivision might take place, all of which must remain compatible with adjacent tetrahedra. The implementation of such a routine thereby becomes lengthy and exceptionally error-prone. For this reason cutting along a surface is planned but not implemented in *Regina* at the present time.

Angle Structures

Angle structures, studied originally by Casson and then developed by Lackenby [22, 23] and Rivin [36, 37], represent a purely algebraic generalisation of hyperbolic structures. An angle structure on an ideal triangulation is formed by assigning an interior dihedral angle to every edge of every tetrahedron in such a way that a variety of linear equations and inequalities are satisfied.

The formation of angle structures is remarkably similar to the formation of normal surfaces, in which a series of triangle and quadrilateral coordinates are assigned to every tetrahedron with a set of linear equations and inequalities similarly imposed upon them. Thus it has been relatively straightforward to extend the code used by *Regina* to enumerate vertex normal surfaces in such a way that the program can also enumerate vertex angle structures.

The inequalities imposed upon an angle structure are that each dihedral angle θ satisfies $0 \leq \theta \leq \pi$. In addition to the enumeration of vertex angle structures, *Regina* can also identify whether a triangulation supports any strict angle structures (structures for which each dihedral angle θ satisfies $0 < \theta < \pi$) or any taut angle structures (structures discussed in [23] for which each dihedral angle is precisely 0 or π).

Splitting Surfaces

Section 4.1 describes splitting surfaces, which when present can offer insight into the 3-manifolds containing them. Section 4.3 in turn defines splitting surface signatures, a compact text-based representation from which a splitting surface and its enclosing 3-manifold can be reconstructed. As well as detecting splitting surfaces as described earlier, *Regina* can reconstruct a 3-manifold from a splitting surface signature and is capable of forming a census of all splitting surface signatures of a given size.

Scripting

Regina offers the ability to write and run arbitrary scripts in *Jython*, a Java implementation of the Python scripting language available from <http://www.jython.org/>. These scripts are essentially high-level programs with immediate access to the mathematical core of *Regina*, and are ideal for performing repetitive tasks over large sets of data. Such tasks might include performing a sequence of tests upon all triangulations in a census, or generating the tables of Appendix B. Scripts can be embedded in *Regina* data files and custom libraries of routines can be written to share code between files.

Interfaces and Documentation

The usual method of running *Regina* provides a full graphical interface that a user can easily understand and use. Alternatively, for those requiring immediate access to the mathematical core of the program, an interactive command-line interface is offered from which users can control the program using the *Jython* scripting language described above. A variety of specialised command-line utilities are also available.

Significant effort has been spent on documentation for the program. A full reference manual is available for end users to assist them in working with *Regina*. For programmers seeking to modify or extend the program or for users writing scripts that interact directly with the mathematical core of *Regina*, the routines offered by both the underlying mathematical engine and the graphical user interface are fully documented.

Data Files

The data files in which a user can save triangulations and other information use a well-organised hierarchical structure. This structure not only allows many triangulations, normal surface lists and other topological structures to be stored together in an organised fashion but also supports the storing of miscellaneous data such as text notes and scripts. The file format is well documented in the reference manual and the files themselves are saved as compressed XML³, allowing simple transfer of native *Regina* data to and from other programs.

1.4 Structure of Material

The structure of this thesis is as follows. Chapters 2 and 3 describe a census of all closed prime minimal 3-manifold triangulations formed from at most six tetrahedra. In Chapter 4 we examine a particular class of normal surfaces called splitting surfaces, and in Chapters 5 and 6 we investigate aspects of the complexity of the normal surface solution space.

A census of 3-manifold triangulations forms a useful reference for testing hypotheses and forming conjectures. The census described in Chapters 2 and 3 includes all closed prime minimal triangulations formed from at most six tetrahedra, including both orientable and non-orientable triangulations. Prime triangulations are those whose underlying 3-manifolds cannot be decom-

³XML is the *Extensible Markup Language*, an open and widely-supported text-based data format.

posed into non-trivial connected sums, and minimal triangulations are those whose underlying 3-manifolds cannot be triangulated using strictly fewer tetrahedra.

The algorithm used in forming the census is covered in detail in Chapter 2. Whilst drawing upon the work of previous authors, a series of original innovations is incorporated to improve the efficiency and generality of the algorithm. These innovations include the use of face pairing automorphisms, testing for 0-efficiency, searching for special subcomplexes within a partially constructed triangulation and the analysis of face pairings.

In Chapter 3 the results of the census thus formed are presented. The majority of this chapter is devoted to developing an understanding of the combinatorial structures of the individual triangulations. Indeed, almost all of the triangulations are categorised into broad families all of whose members share a common large-scale structure and can be analysed simultaneously. Once more this work draws upon and then adds to the results of earlier authors. Original contributions include the description and analysis of chained triangular solid tori, plugged triangular solid tori, square surface bundles and square product pairs. Of particular interest is the non-orientable census, the first such census published to date.

Whereas Chapters 2 and 3 use normal surfaces as a means for obtaining results, Chapter 4 begins a more focused investigation into normal surfaces in their own right. This chapter in particular presents an original study of splitting surfaces, which are normal surfaces satisfying certain structural constraints. Splitting surfaces are of particular interest because they embody the entire combinatorial structure of their enclosing triangulations. Furthermore they offer information regarding the structure of the underlying 3-manifold, and they are straightforward to enumerate for any given triangulation.

In Chapters 5 and 6 we turn to an original investigation of the complexity of the normal surface solution space. Specifically we work towards a bound for the number of maximal embedded faces of the projective solution space. Chapter 5 concentrates on an analysis of edge weight space, a vector space in which we can represent embedded normal surfaces in a particularly efficient manner. Within this chapter the geometric structure of the projective solution space when represented in edge weight space is extensively analysed.

Chapter 6 then examines the particular case in which the projective solution space satisfies a certain set of general position requirements. Here we convert the geometric problem described in Chapter 5 into an algebraic problem, to which a variety of algebraic and combinatorial techniques are applied. The chapter then closes with a discussion of how these results might be generalised.

Finally a series of tables is presented in Appendices A and B, including lists of 3-manifold triangulations and their vertex normal surfaces as found in the census described in Chapters 2 and 3.

Chapter 2

Census of Triangulations

When studying 3-manifold topology, it is frequently the case that one has relatively few examples from which to form conjectures, test hypotheses and draw inspiration. For this reason a census of all 3-manifold triangulations of a specific type is a useful reference.

In this and the following chapter we describe a census of all closed orientable and non-orientable triangulations containing up to six tetrahedra. In these chapters only triangulations with the following properties are of interest.

- *Closed*: The triangulation has no boundary, and in particular has no boundary faces.
- *Real*: The link of each vertex of the triangulation is either a 2-sphere or a disc (and for a closed triangulation must be a 2-sphere). This condition excludes ideal triangulations whose vertex links are tori, Klein bottles and other higher-genus surfaces.
- *Prime*: The underlying 3-manifold of the triangulation cannot be decomposed as a non-trivial connected sum.
- *Minimal*: The underlying 3-manifold of the triangulation cannot be triangulated using strictly fewer tetrahedra.

Note that ideal triangulations are considered not to be closed, since in such triangulations the vertices whose links are not spheres or discs are excluded from the underlying 3-manifold. We nevertheless separate closed and real in the list above, since various components of the census algorithm presented here are easily generalised to include triangulations with only subsets of the above properties.

Both the closed orientable census and the closed non-orientable census as described earlier were created using *Regina*, the computational topology software written by the author and described in Section 1.3. The smaller parts of each census were run on the author's laptop and the more computationally intensive components were run on a series of high-performance computers in the Department of Mathematics and Statistics at the University of Melbourne.

Forming a census of 3-manifold triangulations is not a new idea. In [12] Hildebrand and Weeks form a census of hyperbolic 3-manifolds of up to five tetrahedra and in [5] Callahan, Hildebrand and Weeks extend this census to seven tetrahedra.

Matveev produces a census of closed orientable 3-manifolds of up to six tetrahedra in [29], extended to seven tetrahedra by Ovchinnikov though his results do not appear to be publicly available. In [24] Martelli and Petronio describe a census of closed orientable 3-manifolds of up to nine tetrahedra, the most extensive results known to date.

It is clear then that the census results of these latter authors encompass the closed orientable census of up to six tetrahedra described here. The reasons for presenting the results of this and the following chapter are as follows.

- In addition to an orientable census we also present the results of a census of non-orientable 3-manifold triangulations, described in Section 3.5. Such a census is not included in the works listed above, and many of the techniques used by the above authors are specific to orientable triangulations. The results described in this chapter are in several cases more general.
- The census algorithm presented in this chapter incorporates techniques that are not seen in the previous works. Of particular interest are the use of face pairing automorphisms in Sections 2.2 and 2.3, the use of Jaco and Rubinstein’s 0-efficiency results in Section 2.4.2, the search for special subcomplexes described in Section 2.4.3 and the results regarding face pairing graphs described in Sections 2.6 and 2.7.
- In Chapter 3 the results of the closed orientable census are analysed extensively. The combinatorial structures of the resulting triangulations are examined in detail, and almost all of the triangulations have been categorised into a small number of infinite parameterised families. Such a categorisation, beyond giving more structure to the census results, makes it easier to establish properties of these minimal triangulations and indeed of all triangulations in these infinite families and their underlying 3-manifolds.

Such a categorisation is begun by Matveev in [29] and extended by Martelli and Petronio in [24] and [26]. The material presented in Chapter 3 extends this categorisation, adding to and in some cases generalising the results of these authors. Note that this material was developed independently of Martelli and Petronio’s results.

- The vertex normal surfaces of the minimal triangulations found in each census have been enumerated in Appendix B. As discussed in Section 1.1, the normal surfaces within a triangulation can offer much insight into the structure of the triangulation and of the underlying 3-manifold.

In this chapter we discuss in detail the census algorithm used in *Regina*, and in Chapter 3 we proceed to harvest the results of the orientable and non-orientable censuses thus obtained.

The structure of this chapter is as follows. Section 2.1 describes the large-scale structure of the census algorithm, and the finer details of each algorithm component are discussed from Sections 2.2 to 2.4. In Section 2.5 we measure the performance of this algorithm. Additional results regarding face pairing graphs are proven in Section 2.6, and in Section 2.7 we examine ways in which the algorithm can be made more efficient for future census runs.

2.1 Splitting the Census Algorithm

The census algorithm can be split into two largely independent tasks, these being the generation of face pairings and the generation of corresponding gluing permutations. This is a fairly natural way of approaching a census of triangulations and can be seen back in the earliest hyperbolic census of Hildebrand and Weeks [12].

Algorithm 2.1.16 describes this split more precisely, and Sections 2.2 and 2.3 describe in detail the generation of face pairings and gluing permutations respectively. Before we begin however, it is necessary to define concepts and introduce notation with which we can discuss the various algorithm components.

Notation In a triangulation formed from n tetrahedra, the individual tetrahedra are labelled $\mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_{n-1}$. The four vertices of each tetrahedron \mathcal{T}_t are labelled $\mathcal{V}_{t/0}, \mathcal{V}_{t/1}, \mathcal{V}_{t/2}, \mathcal{V}_{t/3}$. The four faces of each tetrahedron \mathcal{T}_t are labelled $\mathcal{F}_{t/0}, \mathcal{F}_{t/1}, \mathcal{F}_{t/2}, \mathcal{F}_{t/3}$, with face $\mathcal{F}_{t/i}$ opposite vertex $\mathcal{V}_{t/i}$ for each i .

The zero-based notation described above is preferred because, aside from being a natural notation when working in zero-based languages such as C++, it coincides with the way in which triangulations and skeletal information are entered into and presented by *Regina*, the program in which this census algorithm is implemented and with which the resulting triangulations are analysed.

Definitions 2.1.1 (Face Set) For each $n \in \mathbb{N}$, the *face set of order n* , denoted by \mathfrak{F}_n , is defined to be the set of all $4n$ individual tetrahedron faces in an n -tetrahedron triangulation. That is,

$$\mathfrak{F}_n = \{ \mathcal{F}_{t/f} \mid 0 \leq t < n, 0 \leq f < 4 \}.$$

The *augmented face set of order n* , denoted by \mathfrak{F}_n^+ , is defined to be the set

$$\mathfrak{F}_n^+ = \mathfrak{F}_n \cup \{\emptyset\},$$

where \emptyset is referred to as the *boundary* and represents the exterior of the triangulation.

For each triangulation we can then construct a face pairing which details which tetrahedron faces are identified with which others, but without any information regarding which specific vertices are mapped to which by these face identifications.

Definition 2.1.2 (Face Pairing) A *face pairing of order n* is a map $F : \mathfrak{F}_n \rightarrow \mathfrak{F}_n^+$ satisfying the following conditions.

- No $f \in \mathfrak{F}_n$ maps to itself, i.e., $F(f) \neq f$ for all $f \in \mathfrak{F}_n$;
- For any $f, g \in \mathfrak{F}_n$ where $F(f) = g$, it is also true that $F(g) = f$.

Note that the second of these conditions implies that a face pairing is a 1-to-1 map, with the single exception that many elements of \mathfrak{F}_n can map to the boundary \emptyset .

A face pairing of order n can be visually represented as a multigraph on n vertices whose vertices each have degree ≤ 4 ; this representation is discussed in detail in Section 2.6.1. At this

point however we formalise the relationship between an abstract face pairing and the identifications of tetrahedron faces in a triangulation.

Definition 2.1.3 (Associated Face Pairing) For each triangulation T formed from n tetrahedra, the *associated face pairing of T* is the face pairing F of order n defined as follows.

For each tetrahedron face $\mathcal{F}_{t/f}$, if $\mathcal{F}_{t/f}$ is a boundary face of T then $F(\mathcal{F}_{t/f}) = \emptyset$. Otherwise $\mathcal{F}_{t/f}$ must be identified with some other tetrahedron face $\mathcal{F}_{t'/f'}$ of T in which case $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$.

An associated face pairing alone does not provide enough information from which to reconstruct the underlying triangulation. In addition we must know under which rotations and/or reflections the various pairs of faces are identified, i.e., which specific tetrahedron vertices are mapped to which others under each face identification.

Definitions 2.1.4 (Gluing Permutation Selection) Let S_4^+ be defined as $S_4^+ = S_4 \cup \{\emptyset\}$, where S_4 is the set of all permutations on $\{0, 1, 2, 3\}$. If F is a face pairing of order n then a *gluing permutation selection for F* is a map $G : \mathfrak{F}_n \rightarrow S_4^+$ satisfying the following conditions for each $\mathcal{F}_{t/f} \in \mathfrak{F}_n$.

- If $F(\mathcal{F}_{t/f}) = \emptyset$ then $G(\mathcal{F}_{t/f}) = \emptyset$;
- If $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$ for some $\mathcal{F}_{t'/f'} \in \mathfrak{F}_n$ then $G(\mathcal{F}_{t/f})$ is some permutation of S_4 mapping f to f' , and furthermore $G(\mathcal{F}_{t/f})$ and $G(\mathcal{F}_{t'/f'})$ are inverse permutations.

As with face pairings, we can formalise the relationship between an abstract gluing permutation selection and the rotations and/or reflections involved in the face identifications within a triangulation.

Definition 2.1.5 (Associated Gluing Permutation Selection) For each triangulation T , the *associated gluing permutation selection of T* is the gluing permutation selection G defined as follows.

For each tetrahedron face $\mathcal{F}_{t/f}$, if $\mathcal{F}_{t/f}$ is a boundary face of T then $G(\mathcal{F}_{t/f}) = \emptyset$. Otherwise $\mathcal{F}_{t/f}$ must be identified with some other tetrahedron face $\mathcal{F}_{t'/f'}$ in T . Assume that this face identification maps vertices $\mathcal{V}_{t/i}$, $\mathcal{V}_{t/j}$ and $\mathcal{V}_{t/k}$ of tetrahedron \mathcal{T}_t to vertices $\mathcal{V}_{t'/i'}$, $\mathcal{V}_{t'/j'}$ and $\mathcal{V}_{t'/k'}$ of tetrahedron $\mathcal{T}_{t'}$ respectively. Then $\{f, i, j, k\}$ and $\{f', i', j', k'\}$ must each be $\{0, 1, 2, 3\}$ and so we let $G(\mathcal{F}_{t/f})$ be the permutation $p \in S_4$ mapping (f, i, j, k) to (f', i', j', k') .

Lemma 2.1.6 For any triangulation T formed from n tetrahedra, let F be the associated face pairing of T and let G be the associated gluing permutation selection of T . Then F is indeed a face pairing of order n (i.e., F satisfies the conditions of Definition 2.1.2) and G is indeed a gluing permutation selection for F (i.e., G satisfies the conditions of Definition 2.1.4).

Proof These claims are straightforward to verify from Definitions 2.1.3 and 2.1.5. ■

We now verify that an associated face pairing and an associated gluing permutation selection together provide enough information from which to reconstruct the underlying triangulation.

Theorem 2.1.7 Let F be some face pairing and let G be some gluing permutation selection for F . Then there is a unique triangulation T for which F is the associated face pairing of T and G is the associated gluing permutation selection for T .

Proof The number of tetrahedra in T must be the order of F by Definition 2.1.3. We identify the tetrahedron faces of T as follows.

For each face $\mathcal{F}_{t/f}$, if $F(\mathcal{F}_{t/f}) = \emptyset$ then $\mathcal{F}_{t/f}$ must remain a boundary face; otherwise it must be identified with face $F(\mathcal{F}_{t/f})$ as required by Definition 2.1.3. If it is identified with some other face, the corresponding mapping of tetrahedron vertices is determined by permutation $G(\mathcal{F}_{t/f})$ as described in Definition 2.1.5. This determines T completely, so if such a triangulation exists then it must be unique.

Note that each face identification can be seen from two directions, i.e., identifying $\mathcal{F}_{t/f}$ with $\mathcal{F}_{t'/f'}$ according to $F(\mathcal{F}_{t/f})$ and $G(\mathcal{F}_{t/f})$, or alternatively identifying face $\mathcal{F}_{t'/f'}$ with $\mathcal{F}_{t/f}$ according to $F(\mathcal{F}_{t'/f'})$ and $G(\mathcal{F}_{t'/f'})$. The conditions imposed upon a face pairing and a gluing permutation selection in Definitions 2.1.2 and 2.1.4 respectively ensure that no inconsistencies will be encountered as a result of these two approaches, so that they must in fact describe the same face identification. Thus T can indeed be constructed as described.

Finally a simple check verifies that F and G are the associated face pairing and associated gluing permutation selection for T as required. ■

Note that in Theorem 2.1.7 we make no claims as to whether triangulation T actually represents a 3-manifold. The vertices of T for instance might have links that are neither spheres nor discs, and the edges of T might be identified with themselves in reverse causing their midpoints to have links that are projective planes. The triangulation might even be disconnected, consisting of several smaller triangulations each using a subset of the available tetrahedra.

Definition 2.1.8 (Connectedness) A face pairing F of order n is *connected* if for any pair of tetrahedra \mathcal{T}_a and \mathcal{T}_b ($0 \leq a, b < n$) there is some sequence of tetrahedra

$$\mathcal{T}_a = \mathcal{T}_{i_0}, \mathcal{T}_{i_1}, \mathcal{T}_{i_2}, \dots, \mathcal{T}_{i_k} = \mathcal{T}_b$$

in which each two consecutive tetrahedra are joined by F . That is, for each $j = 0, 1, \dots, k-1$, there is some face $\mathcal{F}_{i_j/f}$ of tetrahedron \mathcal{T}_{i_j} for which $F(\mathcal{F}_{i_j/f})$ is a face of tetrahedron $\mathcal{T}_{i_{j+1}}$.

Lemma 2.1.9 A triangulation is connected if and only if its associated face pairing is connected.

Proof This result is immediate from Definition 2.1.8, since a triangulation formed from n tetrahedra is connected if and only if for any pair of tetrahedra \mathcal{T}_a and \mathcal{T}_b in the triangulation there is some sequence of tetrahedra beginning with \mathcal{T}_a and ending with \mathcal{T}_b in which each two consecutive tetrahedra are adjacent. ■

From Theorem 2.1.7 then we can anticipate the census algorithm by observing that we can split this algorithm into two components as follows.

1. Construct all face pairings of a given order. Connectedness can be enforced at this stage, as can constraints relating to the number of boundary faces. For instance, if we only seek closed triangulations then we can require that no constructed face pairing maps any face to \emptyset .

2. For each face pairing F that has been constructed, generate all gluing permutation selections G for F . The unique triangulation associated with F and G can then be examined as to whether it indeed represents a 3-manifold and whether it is a desirable member of the census.

In practice, the process in which gluing permutations are generated can be pruned so that many triangulations that we know will not be desirable members of the census are in fact never generated at all, resulting in a dramatic improvement in the total running time. This pruning is detailed in Section 2.3.1.

Furthermore, most triangulations will be created many times by the above process according to the different orders in which their tetrahedra and the vertices within these tetrahedra can be labelled. In order to keep the number of generated triangulations to a minimum we will ensure that no triangulation is constructed more than once up to relabelling.

Definitions 2.1.10 (Relabelling) A *relabelling* of order n is a permutation ρ on \mathfrak{F}_n for which the following condition is satisfied.

- For each t , f_1 and f_2 , if $\rho(\mathcal{F}_{t/f_1}) = \mathcal{F}_{t'_1/f'_1}$ and $\rho(\mathcal{F}_{t/f_2}) = \mathcal{F}_{t'_2/f'_2}$ then $t'_1 = t'_2$. That is, any two faces of the same tetrahedron \mathcal{T}_t must be mapped under ρ to two faces of the same tetrahedron $\mathcal{T}_{t'}$.

For convenience we sometimes extend ρ to be a permutation on \mathfrak{F}_n^+ by defining $\rho(\emptyset) = \emptyset$.

In effect then, a relabelling corresponds to a renumbering of the n tetrahedra of a triangulation followed by renumberings of the four faces (or equivalently of the four vertices) of each individual tetrahedron. So if T is a triangulation formed from n tetrahedra, let $\rho(T)$ be the triangulation obtained from T by renumbering the tetrahedra and their individual faces and vertices as described.

If F is a face pairing of order n , let $\rho(F)$ be the face pairing F' similarly obtained from F by this relabelling, i.e., for which $F'(\rho(\mathcal{F}_{t/f})) = \rho(F(\mathcal{F}_{t/f}))$. Note that this yields the relation

$$\rho(F)(\mathcal{F}_{t/f}) = \rho(F(\rho^{-1}(\mathcal{F}_{t/f}))) \quad (2.1)$$

which is used frequently throughout this chapter.

If G is a gluing permutation selection for F , let $\rho(G)$ be the gluing permutation selection G' defined as follows. For each t and f , if $G(\mathcal{F}_{t/f}) = \emptyset$ then $G'(\rho(\mathcal{F}_{t/f})) = \emptyset$. Otherwise $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$ for some t', f' and $G(\mathcal{F}_{t/f}) = p$ for some $p \in S_4$. Let $\pi \in S_4$ be the permutation for which ρ maps $\mathcal{F}_{t/z}$ to $\mathcal{F}_{s/\pi(z)}$ for some s and for each $z = 0, 1, 2, 3$. Let $\pi' \in S_4$ be the permutation for which ρ maps $\mathcal{F}_{t'/z}$ to $\mathcal{F}_{s'/\pi'(z)}$ for some s' and for each $z = 0, 1, 2, 3$. Then $G'(\rho(\mathcal{F}_{t/f})) = \pi' \circ p \circ \pi^{-1}$.

Finally from the condition above we see that for any tetrahedron \mathcal{T}_t , all four faces of \mathcal{T}_t are mapped under ρ to the four faces of some single tetrahedron $\mathcal{T}_{t'}$. We thus define $\rho(\mathcal{T}_t) = \mathcal{T}_{t'}$, and since ρ is a permutation on \mathfrak{F}_n we observe that the map $\mathcal{T}_t \mapsto \rho(\mathcal{T}_t)$ is likewise a permutation on the set of tetrahedra $\{\mathcal{T}_0, \dots, \mathcal{T}_{n-1}\}$.

Lemma 2.1.11 Let ρ be a relabelling of order n . If F is a face pairing of order n then $\rho(F)$ also satisfies the requirements of a face pairing of order n . If G is a gluing permutation selection for F then $\rho(G)$ satisfies the requirements of a gluing permutation selection for $\rho(F)$. Finally, if F and G are the associated face pairing and associated gluing permutation selection for triangulation T

then $\rho(F)$ and $\rho(G)$ are the associated face pairing and associated gluing permutation selection for triangulation $\rho(T)$.

Proof For any face pairing F it is simple to construct a gluing permutation selection G for F as follows. If $F(\mathcal{F}_{t/f}) = \emptyset$ then we set $G(\mathcal{F}_{t/f}) = \emptyset$. Otherwise for each pair $\mathcal{F}_{t/f}$ and $\mathcal{F}_{t'/f'}$ with $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$ and $F(\mathcal{F}_{t'/f'}) = \mathcal{F}_{t/f}$ we let $G(\mathcal{F}_{t/f}) = \pi$ and $G(\mathcal{F}_{t'/f'}) = \pi^{-1}$ for some arbitrary permutation $\pi \in S_4$ mapping f to f' .

Similarly, for any face pairing F and any gluing permutation selection G for F , Theorem 2.1.7 shows that there is some triangulation T for which F is the associated face pairing and G is the associated gluing permutation selection.

We thus prove only the final statement of this lemma, i.e., that if F and G are the associated face pairing and gluing permutation selection for some triangulation T then $\rho(F)$ and $\rho(G)$ are the associated face pairing and gluing permutation selection for triangulation $\rho(T)$. The remaining statements of this lemma then follow immediately from Lemma 2.1.6.

Consider then triangulation T with associated face pairing F and associated gluing permutation selection G . Let $F' = \rho(F)$ and $G' = \rho(G)$. Consider any tetrahedron face in \mathfrak{F}_n ; since ρ is a permutation this face can be written as $\rho(\mathcal{F}_{t/f})$ for some $\mathcal{F}_{t/f} \in \mathfrak{F}_n$.

If $\rho(\mathcal{F}_{t/f})$ is a boundary face of $\rho(T)$ then $\mathcal{F}_{t/f}$ is a boundary face of T and so $F(\mathcal{F}_{t/f}) = \emptyset$. Hence $F'(\rho(\mathcal{F}_{t/f})) = G'(\rho(\mathcal{F}_{t/f})) = \emptyset$ by Definitions 2.1.10, as required for an associated face pairing or associated gluing permutation selection for $\rho(T)$.

Otherwise $\rho(\mathcal{F}_{t/f})$ is identified with some other face $\rho(\mathcal{F}_{t'/f'})$ of $\rho(T)$ in which case $\mathcal{F}_{t/f}$ is identified with $\mathcal{F}_{t'/f'}$ in T . This implies that $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$ and so $F'(\rho(\mathcal{F}_{t/f})) = \rho(F(\mathcal{F}_{t/f})) = \rho(\mathcal{F}_{t'/f'})$, completing the requirements for F' to be the associated face pairing of triangulation $\rho(T)$.

Furthermore, suppose the identification of $\mathcal{F}_{t/f}$ and $\mathcal{F}_{t'/f'}$ in T maps vertices $\mathcal{V}_{t/i}$, $\mathcal{V}_{t/j}$ and $\mathcal{V}_{t/k}$ of tetrahedron \mathcal{T}_t to vertices $\mathcal{V}_{t'/i'}$, $\mathcal{V}_{t'/j'}$ and $\mathcal{V}_{t'/k'}$ of tetrahedron $\mathcal{T}_{t'}$. Let $p \in S_4$ be the permutation mapping (f, i, j, k) to (f', i', j', k') , i.e., let $p = G(\mathcal{F}_{t/f})$. As in Definitions 2.1.10, let $\pi, \pi' \in S_4$ be the permutations for which ρ maps $\mathcal{F}_{t/z}$ to $\mathcal{F}_{s/\pi(z)}$ and $\mathcal{F}_{t'/z}$ to $\mathcal{F}_{s'/\pi'(z)}$ for some s, s' and for each $z = 0, 1, 2, 3$.

Then if the identification of faces $\mathcal{F}_{t/f}$ and $\mathcal{F}_{t'/f'}$ maps vertex $\mathcal{V}_{t/z}$ to vertex $\mathcal{V}_{t'/z'}$ in triangulation T , it follows that the identification of faces $\rho(\mathcal{F}_{t/f})$ and $\rho(\mathcal{F}_{t'/f'})$ maps vertex $\mathcal{V}_{s/\pi(z)}$ to vertex $\mathcal{V}_{s'/\pi'(z')}$ in triangulation $\rho(T)$. Since $z' = p(z)$ we see that $\pi'(z') = (\pi' \circ p \circ \pi^{-1})(\pi(z))$. Thus $G'(\rho(\mathcal{F}_{t/f})) = \pi' \circ p \circ \pi^{-1}$ is indeed the permutation mapping $(\pi(f), \pi(i), \pi(j), \pi(k))$ to $(\pi'(f'), \pi'(i'), \pi'(j'), \pi'(k'))$, completing the requirements for G' to be the associated gluing permutation selection for triangulation $\rho(T)$. ■

Since the census algorithm is split into the separate processes of generating face pairings and generating gluing permutation selections, we investigate how the problem of avoiding relabellings of the same triangulation can be likewise split between these two processes.

Definitions 2.1.12 (Isomorphism Class) For any triangulation or face pairing x , the *isomorphism class* of x is the set of all triangulations or face pairings that can be obtained from x by a relabelling. Any two triangulations or face pairings in the same isomorphism class are said to be *isomorphic*.

Definitions 2.1.13 (*F*-Isomorphism Class) Let F be any face pairing. For any gluing permutation selection x for F , the *F-isomorphism class* of x is the set of all gluing permutation selections for F that can be obtained from x by a relabelling ρ for which $\rho(F) = F$. Any two gluing permutation selections in the same *F-isomorphism class* are said to be *F-isomorphic*.

Lemma 2.1.14 Let F and F' be isomorphic face pairings. Then any triangulation with associated face pairing F is isomorphic to some triangulation with associated face pairing F' .

Likewise, let G and G' be *F-isomorphic* gluing permutation selections for some face pairing F . Then any triangulation with associated face pairing F and gluing permutation G is isomorphic to some triangulation with associated face pairing F and gluing permutation G' .

Proof For the first claim let ρ be a relabelling for which $\rho(F) = F'$. Then if triangulation T has associated face pairing F it follows from Lemma 2.1.11 that triangulation $\rho(T)$ has associated face pairing F' .

For the second claim let ρ be a relabelling for which $\rho(F) = F$ and $\rho(G) = G'$. Then if triangulation T has associated face pairing F and gluing permutation selection G it follows again from Lemma 2.1.11 that triangulation $\rho(T)$ has associated face pairing F and gluing permutation selection G' . ■

From Lemma 2.1.14 it follows that in our census algorithm we need only generate one face pairing from each isomorphism class, and for each such face pairing F we need only generate one gluing permutation selection from each *F-isomorphism class*.

In fact this is sufficient to ensure that each triangulation is produced once and only once. Lemma 2.1.11 shows that all relabellings of a triangulation have isomorphic associated face pairings, and if two relabellings of a triangulation have identical associated face pairings F then these triangulations have *F-isomorphic* associated gluing permutations. If two relabellings of a triangulation have identical associated face pairings and identical associated gluing permutations then Theorem 2.1.7 shows that they are in fact the same triangulation.

Before presenting the final structure of the census algorithm, we pause to examine a tool that will help improve the overall running time of the census.

Definition 2.1.15 (Automorphism) For any triangulation or face pairing x , an *automorphism* of x is a relabelling ρ for which $\rho(x) = x$.

By generating automorphisms of face pairings we can reduce the running time for subsequent parts of the census algorithm, as described in Sections 2.2 and 2.3. The overall structure of our census algorithm is thus as follows.

Algorithm 2.1.16 (Census of Triangulations) To form a census of all connected 3-manifold triangulations formed from n tetrahedra and satisfying some set of constraints, in which each triangulation is generated precisely once up to isomorphism, we perform the following steps.

1. Generate one representative from each possible isomorphism class of face pairings of order n , as well as the set of all automorphisms for each such face pairing. Face pairings that are not connected or that can never lead to triangulations satisfying the given constraints (such as pairings with faces mapped to \emptyset in a census of closed triangulations) need not be generated.

2. For each generated face pairing F , generate one representative from each possible F -isomorphism class of gluing permutation selections for F . Again, gluing permutation selections that cannot lead to triangulations satisfying the given constraints (such as those representing non-orientable triangulations in a census of orientable triangulations) need not be generated.
3. For each generated gluing permutation selection G , construct the unique triangulation associated with F and G , check whether it satisfies the full set of census constraints, and if so then include it in the final census.

The first two of these steps will be examined in detail in Sections 2.2 and 2.3. The final step is dependent upon the specific census constraints. For the constraints described at the beginning of this chapter we discuss this final step in Section 2.4.

2.2 Generating Face Pairings

In this section we examine the details of step 1 of Algorithm 2.1.16, namely the generation of face pairings. The simplest way to ensure we have precisely one representative from each isomorphism class of face pairings is to impose a total order on face pairings, generate all possible face pairings and then throw away any face pairing F for which there exists a relabelling ρ with $\rho(F) < F$. The face pairings that remain will be the unique minimal representatives of the various face pairing isomorphism classes.

Definition 2.2.1 (Ordering of Face Pairings) Let F_1 and F_2 be face pairings of order n . We will declare that $F_1 < F_2$ if and only if the sequence

$$F_1(\mathcal{F}_{0/0}), F_1(\mathcal{F}_{0/1}), F_1(\mathcal{F}_{0/2}), F_1(\mathcal{F}_{0/3}), F_1(\mathcal{F}_{1/0}), \dots, F_1(\mathcal{F}_{n-1/3})$$

is lexicographically smaller than the sequence

$$F_2(\mathcal{F}_{0/0}), F_2(\mathcal{F}_{0/1}), F_2(\mathcal{F}_{0/2}), F_2(\mathcal{F}_{0/3}), F_2(\mathcal{F}_{1/0}), \dots, F_2(\mathcal{F}_{n-1/3}),$$

where the individual elements of \mathfrak{F}_n^+ are ordered as

$$\mathcal{F}_{0/0} < \mathcal{F}_{0/1} < \mathcal{F}_{0/2} < \mathcal{F}_{0/3} < \mathcal{F}_{1/0} < \dots < \mathcal{F}_{n-1/3} < \emptyset.$$

There is an immediate efficiency concern however with the algorithm described above. Since we only want one representative from what could be up to $(4!)^n n!$ members of each isomorphism class, it is terribly inefficient to generate every possible face pairing. Instead we can impose constraints upon the face pairing generation so that a large number of face pairings that are not minimal representatives of their isomorphism classes are never generated at all.

With this aim in mind, the generation of face pairings in *Regina* is done as follows.

Algorithm 2.2.2 (Face Pairing Generation) The following procedure constructs every connected face pairing F of order n that is a minimal representative of its isomorphism class.

Recursively select values for $F(\mathcal{F}_{0/0}), F(\mathcal{F}_{0/1}), \dots, F(\mathcal{F}_{n-1/3})$ in turn. If some $F(\mathcal{F}_{t/f})$ is already determined, i.e. if there is some $\mathcal{F}_{t'/f'} < \mathcal{F}_{t/f}$ for which $F(\mathcal{F}_{t'/f'}) = \mathcal{F}_{t/f}$ and hence $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$, we skip over face $\mathcal{F}_{t/f}$ and move immediately on to the following face.

At each stage, when selecting $F(\mathcal{F}_{t/f})$ we try all possible values subject to the constraints below. For each value that satisfies these constraints we recursively select the image under F of the following face as described above, and once we run out of candidate values for $F(\mathcal{F}_{t/f})$ we backtrack and try the next possible image under F of the face prior to $\mathcal{F}_{t/f}$.

When selecting possible values for $F(\mathcal{F}_{t/f})$, the following constraints are imposed in order to avoid generating face pairings that either cannot be minimal representatives of their isomorphism classes or cannot lead to a triangulation that satisfies the census constraints. Note that these constraints allow us to avoid constructing many but not all such undesirable face pairings.

- Only try $F(\mathcal{F}_{t/f}) > \mathcal{F}_{t/f}$, since otherwise $F(\mathcal{F}_{t/f})$ would have already been determined at an earlier stage of the recursion.
- Do not try $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$ if no face is yet mapped to $\mathcal{F}_{t'-1/i}$ for any i , since otherwise a relabelling that simply switches tetrahedra t' and $t' - 1$ will lead to a smaller face pairing.
- Do not try $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$ if no face is yet mapped to $\mathcal{F}_{t'/f'-1}$, since otherwise a relabelling that simply switches $\mathcal{F}_{t'/f'}$ and $\mathcal{F}_{t'/f'-1}$ will lead to a smaller face pairing.
- Do not try $F(\mathcal{F}_{t/f}) < F(\mathcal{F}_{t/f'})$ for any $f' < f$, since this implies $F(\mathcal{F}_{t/f'}) > F(\mathcal{F}_{t/f}) > \mathcal{F}_{t/f} > \mathcal{F}_{t/f'}$ and so a relabelling that simply switches $\mathcal{F}_{t/f}$ and $\mathcal{F}_{t/f'}$ will lead to a smaller face pairing.
- If $F(\mathcal{F}_{0/0}) \neq \mathcal{F}_{0/1}$ and $t > 0$, do not try $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t/f'}$ for any f' , since otherwise a relabelling that switches tetrahedra 0 and t in a way that switches faces $\mathcal{F}_{t/f}$ and $\mathcal{F}_{t/f'}$ with faces $\mathcal{F}_{0/0}$ and $\mathcal{F}_{0/1}$ respectively will lead to a smaller face pairing.
- Do not close off a subset of tetrahedra. Specifically, if we are selecting $F(\mathcal{F}_{t/0})$ for some $t > 0$ then we must backtrack immediately, since from the constraints above we see that for each face $\mathcal{F}_{t'/f'} < \mathcal{F}_{t/0}$ it is true that $F(\mathcal{F}_{t'/f'}) < \mathcal{F}_{t/0}$ and so the resulting face pairing will not be connected.

Finally, for each complete face pairing F that is generated, we determine whether F is the minimal representative of its isomorphism class and if so we generate a full list of its automorphisms. This procedure is described separately in Algorithm 2.2.5.

Determining whether F is a minimal representative of its isomorphism class is a non-trivial task. The basic idea is to run through all possible relabellings. For each relabelling ρ , if $\rho(F) < F$ then F is not a minimal representative, and if $\rho(F) = F$ then we can add ρ to the list of automorphisms of F .

Unfortunately this procedure as presented is infeasible since there are $(4!)^n n!$ relabellings of order n . We therefore modify this algorithm so that it is only necessary to examine a much smaller number set of relabellings.

Definition 2.2.3 (*F*-Minimal Relabelling) For each face pairing F of order n , an *F*-minimal relabelling is a relabelling ρ of order n for which there is no other relabelling ρ' of order n satisfying $\rho'(F) < \rho(F)$.

Lemma 2.2.4 Let F be a face pairing. Then every F -minimal relabelling ρ satisfies $\rho(F) \leq F$. Furthermore, F is a minimal representative of its isomorphism class if and only if there is no F -minimal relabelling ρ for which $\rho(F) < F$. If F is indeed a minimal representative of its isomorphism class then every automorphism of F is also F -minimal.

Proof Note first that if $\rho(F) > F$ for some relabelling ρ then the identity relabelling ι satisfies $\iota(F) < \rho(F)$ and so ρ cannot be F -minimal.

If F is a minimal representative of its isomorphism class then clearly there is no F -minimal relabelling ρ for which $\rho(F) < F$. On the other hand, say F is not a minimal representative of its isomorphism class. Then there is some relabelling ρ for which $\rho(F) < F$. If ρ is not F -minimal, there is some ρ' with $\rho'(F) < \rho(F) < F$. If ρ' is also not F -minimal then there is some ρ'' with $\rho''(F) < \rho'(F)$ and so on, continuing down a chain of smaller and smaller relabellings until we reach some F -minimal relabelling ρ_0 with $\rho_0(F) < \dots < \rho(F) < F$. Note that this procedure must terminate since there are only finitely many relabellings of any given order.

Finally, let F be a minimal representative of its isomorphism class and let ρ be one of its automorphisms. If ρ is not F -minimal then there is some relabelling ρ' for which $\rho'(F) < \rho(F)$. Since ρ is an automorphism however this implies that $\rho'(F) < F$, contradicting the minimality of F . ■

From Lemma 2.2.4 we see that our algorithm for determining whether a face pairing F is a minimal representative of its isomorphism class can be modified so that we construct as few non- F -minimal relabellings as possible. In addition we will construct our relabellings in such a way that if $\rho(F) < F$ then this fact is discovered quickly, since such a discovery allows us to immediately terminate the algorithm.

Algorithm 2.2.5 (Face Pairing Minimality) Given a face pairing F , we wish to determine whether F is a minimal representative of its isomorphism class and, if so, to construct a full list of the automorphisms of F . Both of these tasks involve searching for relabellings ρ for which either $\rho(F) < F$ (implying that F is non-minimal) or $\rho(F) = F$ (implying that ρ is an automorphism of F). In particular all F -minimal relabellings are examined.

We construct each relabelling ρ by recursively selecting a preimage for each face in the order $\rho^{-1}(\mathcal{F}_{0/0}), \rho^{-1}(\mathcal{F}_{0/1}), \dots, \rho^{-1}(\mathcal{F}_{n-1/3})$. As an exception to this ordering, whenever $\rho^{-1}(\mathcal{F}_{t/f})$ is selected and $F(\mathcal{F}_{t/f}) \neq \emptyset$ then the preimage $\rho^{-1}(F(\mathcal{F}_{t/f}))$ can be determined immediately; this is described in detail further on. It will be observed that as this recursion progresses the following invariant will remain true.

Invariant: At each stage of the algorithm, whilst selecting the preimage $\rho^{-1}(\mathcal{F}_{t/f})$ we are guaranteed that ρ^{-1} as far as it has already been constructed satisfies the requirements of an automorphism. That is, for each face $\mathcal{F}_{t'/f'} < \mathcal{F}_{t/f}$ we are guaranteed that $\rho^{-1}(F(\mathcal{F}_{t'/f'}))$ has already been determined and that $F(\rho^{-1}(\mathcal{F}_{t'/f'})) = \rho^{-1}(F(\mathcal{F}_{t'/f'}))$, where we let $\rho^{-1}(\emptyset) = \emptyset$ for convenience.

As the preimages $\rho^{-1}(\mathcal{F}_{0/0})$, $\rho^{-1}(\mathcal{F}_{0/1})$, \dots , $\rho^{-1}(\mathcal{F}_{n-1/3})$ are selected in turn, each preimage will be chosen from a restricted set of possibilities as follows.

- When selecting the initial preimage $\rho^{-1}(\mathcal{F}_{0/0})$, all $4n$ possible values for $\rho^{-1}(\mathcal{F}_{0/0}) \in \mathfrak{F}_n$ are attempted.

We make an exception however when $F(\mathcal{F}_{0/0}) = \mathcal{F}_{0/1}$, which from the constraints imposed in Algorithm 2.2.2 can be seen to occur precisely when there is some tetrahedron two of whose faces are to be identified. In this case we only try preimages $\rho^{-1}(\mathcal{F}_{0/0})$ for which $\rho^{-1}(\mathcal{F}_{0/0})$ and $F(\rho^{-1}(\mathcal{F}_{0/0}))$ belong to the same tetrahedron.

The reason for this restriction is as follows. Say $\rho^{-1}(\mathcal{F}_{0/0})$ and $F(\rho^{-1}(\mathcal{F}_{0/0}))$ belong to different tetrahedra. This in turn implies that $\rho(\rho^{-1}(\mathcal{F}_{0/0}))$ and $\rho(F(\rho^{-1}(\mathcal{F}_{0/0})))$ belong to different tetrahedra. Since $\rho(\rho^{-1}(\mathcal{F}_{0/0})) = \mathcal{F}_{0/0}$ and $\rho(F(\rho^{-1}(\mathcal{F}_{0/0}))) = \rho(F)(\mathcal{F}_{0/0})$ we conclude that $\rho(F)(\mathcal{F}_{0/0})$ does not belong to tetrahedron \mathcal{T}_0 and so $\rho(F)(\mathcal{F}_{0/0}) > \mathcal{F}_{0/1} = F(\mathcal{F}_{0/0})$. Thus $\rho(F) > F$ and ρ cannot be F -minimal.

- When selecting each remaining preimage $\rho^{-1}(\mathcal{F}_{t/f})$, the tetrahedron of $\rho^{-1}(\mathcal{F}_{t/f})$ is already predetermined as follows. If $f > 0$ then the tetrahedron must be the same as for $\rho^{-1}(\mathcal{F}_{t/0})$ according to Definitions 2.1.10. If $f = 0$ then the connectedness constraint of Algorithm 2.2.2 implies that $F(\mathcal{F}_{t/f}) < \mathcal{F}_{t/f}$. In this case the preimage $\rho^{-1}(\mathcal{F}_{t/f}) = \rho^{-1}(F(F(\mathcal{F}_{t/f})))$ was already determined when $\rho^{-1}(F(\mathcal{F}_{t/f}))$ was chosen at an earlier stage of the algorithm.

In either case, since the tetrahedron for $\rho^{-1}(\mathcal{F}_{t/f})$ is already determined, only the faces of that tetrahedron that do not already have images under ρ (of which there are at most three) need be considered as possible values for $\rho^{-1}(\mathcal{F}_{t/f})$.

Assume that we have just selected the preimage $\rho^{-1}(\mathcal{F}_{t/f})$; let this preimage be $\rho^{-1}(\mathcal{F}_{t/f}) = \mathcal{F}_{t_0/f_0}$. Consider then the partial sequence $\rho(F)(\mathcal{F}_{0/0})$, $\rho(F)(\mathcal{F}_{0/1})$, \dots , $\rho(F)(\mathcal{F}_{t/f})$. From Definition 2.2.1 we see that if ρ is to have any hope of becoming an F -minimal relabelling then this partial sequence must be as lexicographically small as possible.

From Equation 2.1 we see that each term $\rho(F)(\mathcal{F}_{t'/f'})$ of this sequence can be expressed as $\rho(F(\rho^{-1}(\mathcal{F}_{t'/f'})))$ which from our invariant is already known to be $\rho(\rho^{-1}(F(\mathcal{F}_{t'/f'}))) = F(\mathcal{F}_{t'/f'})$ for each $\mathcal{F}_{t'/f'} < \mathcal{F}_{t/f}$. Thus the only term of this sequence in which we might still have any choice is the final term $\rho(F)(\mathcal{F}_{t/f}) = \rho(F(\rho^{-1}(\mathcal{F}_{t/f})))$, which must therefore be minimised. We thus temporarily break the order of construction of ρ to determine $\rho(F(\rho^{-1}(\mathcal{F}_{t/f}))) = \rho(F(\mathcal{F}_{t_0/f_0}))$ as follows.

- If $F(\mathcal{F}_{t_0/f_0}) = \emptyset$ or $\rho(F(\mathcal{F}_{t_0/f_0}))$ has already been determined then there is no choice to be made.
- Otherwise let $F(\mathcal{F}_{t_0/f_0}) = \mathcal{F}_{t_1/f_1}$. If some face of tetrahedron \mathcal{T}_{t_1} has already been assigned an image under ρ then the tetrahedron $\rho(\mathcal{T}_{t_1})$ has already been established. In this case we choose $\rho(F(\mathcal{F}_{t_0/f_0}))$ to be the first unused face of tetrahedron $\rho(\mathcal{T}_{t_1})$.

If no face of tetrahedron \mathcal{T}_{t_1} has yet been assigned an image under ρ then let tetrahedron \mathcal{T}_{t_2} be the first tetrahedron for which no faces have yet been assigned a preimage under ρ . We then choose $\rho(F(\mathcal{F}_{t_0/f_0})) = \rho(\mathcal{F}_{t_1/f_1})$ to be the face $\mathcal{F}_{t_2/0}$.

By selecting $\rho(F(\mathcal{F}_{t_0/f_0}))$ using this method we ensure that the partial sequence described earlier is made as small as possible. In particular, any other selection for $\rho(F(\mathcal{F}_{t_0/f_0}))$ will result in a relabelling ρ that is not F -minimal and can therefore be ignored.

Now that $\rho(F(\rho^{-1}(\mathcal{F}_{t/f})))$ has been defined we can compare the relative order of face pairings $\rho(F)$ and F as far as they have been constructed. Specifically we lexicographically compare the partial sequences $\rho(F)(\mathcal{F}_{0/0}), \rho(F)(\mathcal{F}_{0/1}), \dots, \rho(F)(\mathcal{F}_{t/f})$ and $F(\mathcal{F}_{0/0}), F(\mathcal{F}_{1/1}), \dots, F(\mathcal{F}_{t/f})$.

As observed earlier, our invariant ensures that $\rho(F)(\mathcal{F}_{t'/f'}) = F(\mathcal{F}_{t'/f'})$ for each face $\mathcal{F}_{t'/f'} < \mathcal{F}_{t/f}$. Thus the relative order of these two sequences depends entirely upon the relative order of the final terms $\rho(F)(\mathcal{F}_{t/f})$ and $F(\mathcal{F}_{t/f})$. Recall from Equation 2.1 that $\rho(F)(\mathcal{F}_{t/f}) = \rho(F(\rho^{-1}(\mathcal{F}_{t/f})))$ which we have recently determined as described above.

- If $\rho(F(\rho^{-1}(\mathcal{F}_{t/f}))) < F(\mathcal{F}_{t/f})$ then we will have $\rho(F) < F$. In this case F is not a minimal representative of its isomorphism class and we can terminate the algorithm.
- If $\rho(F(\rho^{-1}(\mathcal{F}_{t/f}))) > F(\mathcal{F}_{t/f})$ then we will have $\rho(F) > F$. In this case ρ will not be an F -minimal relabelling. We therefore abandon the currently selected value for $\rho^{-1}(\mathcal{F}_{t/f})$ and try a new one.
- If $\rho(F(\rho^{-1}(\mathcal{F}_{t/f}))) = F(\mathcal{F}_{t/f})$ then we cannot yet determine the relative order of $\rho(F)$ and F . We do however observe that $F(\rho^{-1}(\mathcal{F}_{t/f})) = \rho^{-1}(F(\mathcal{F}_{t/f}))$ and so our invariant, which states that ρ^{-1} is a partial automorphism, has been maintained.

Note also that in the process of selecting the preimage $\rho^{-1}(\mathcal{F}_{t/f})$ we have determined the partner preimage $\rho^{-1}(F(\mathcal{F}_{t/f}))$ as promised earlier. At this point we return to our recursive selection of preimages and choose a preimage for the face following $\mathcal{F}_{t/f}$.

If there are no more preimages to select then we must have a complete relabelling ρ for which $\rho(F) = F$. In this case we add ρ to our list of automorphisms of F .

Once our search for relabellings is complete, if we did not terminate with $\rho(F) < F$ then Lemma 2.2.4 tells us that F is indeed a minimal representative of its isomorphism class. Furthermore this same lemma shows that the list of automorphisms that was constructed does indeed contain every automorphism of F .

It is worth noting that, in the overall running time when forming a census, the time taken to generate face pairings is utterly negligible. For the 8 tetrahedron closed orientable census which consumed 27 months of processor time as seen in Table 2.1 on page 65, the generation of face pairings took just 5 seconds.

2.3 Generating Gluing Permutations

We now proceed to step 2 of Algorithm 2.1.16 which involves the generation of gluing permutation selections for some face pairing F . Recall from Algorithm 2.1.16 that we require precisely one gluing permutation selection from each F -isomorphism class.

As with the generation of face pairings in Section 2.2, this is done by imposing a total order on gluing permutation selections for F . We generate all possible gluing permutation selections

for F and discard all gluing permutation selections G for which there is some relabelling ρ with $\rho(G) < G$ and $\rho(F) = F$. The gluing permutation selections that survive this process will be the unique minimal representatives of the various F -isomorphism classes.

Definition 2.3.1 (Ordering of Gluing Permutation Selections) We impose a total order on the augmented set of permutations $S_4^+ = S_4 \cup \{\emptyset\}$ as follows. Let $p_1, p_2 \in S_4$ be permutations on $\{0, 1, 2, 3\}$. We declare that $p_1 < p_2$ if and only if the sequence $p_1(0), \dots, p_1(3)$ is lexicographically smaller than the sequence $p_2(0), \dots, p_2(3)$. Finally we declare that \emptyset is larger than any permutation in S_4 .

Now let G_1 and G_2 be gluing permutation selections for some face pairing F . We declare that $G_1 < G_2$ if and only if the sequence

$$G_1(\mathcal{F}_{0/0}), G_1(\mathcal{F}_{0/1}), G_1(\mathcal{F}_{0/2}), G_1(\mathcal{F}_{0/3}), G_1(\mathcal{F}_{1/0}), \dots, G_1(\mathcal{F}_{n-1/3})$$

is lexicographically smaller than the sequence

$$G_2(\mathcal{F}_{0/0}), G_2(\mathcal{F}_{0/1}), G_2(\mathcal{F}_{0/2}), G_2(\mathcal{F}_{0/3}), G_2(\mathcal{F}_{1/0}), \dots, G_2(\mathcal{F}_{n-1/3}).$$

As with the generation of face pairings in Section 2.2, we construct gluing permutation selections G by recursively selecting $G(\mathcal{F}_{t/f})$ for each tetrahedron face $\mathcal{F}_{t/f}$. Unlike the generation of face pairings however, we cannot enforce as many constraints during this recursion to eliminate unwanted relabellings of the same gluing permutation selection.

If only prime minimal triangulations are required in the census then there are constraints that can be imposed to eliminate many triangulations that do not fit these criteria. These additional constraints are detailed in Section 2.3.1.

Algorithm 2.3.2 (Gluing Permutation Generation) Assume we are given a face pairing F and a list of its automorphisms as produced by Algorithm 2.2.2. The following procedure constructs a variety of gluing permutation selections for F each of which is a minimal representative of its F -isomorphism class. For any F -isomorphism class from which no gluing permutation is constructed, it is guaranteed that the corresponding triangulations do not satisfy the census constraints.

To construct gluing permutation selection G , we recursively select the individual permutations $G(\mathcal{F}_{0/0}), G(\mathcal{F}_{0/1}), \dots, G(\mathcal{F}_{n-1/3})$ in turn. Each time we select permutation $G(\mathcal{F}_{t/f})$, we must also declare permutation $G(F(\mathcal{F}_{t/f}))$ to be the inverse permutation as required by Definition 2.1.4. At each stage when selecting $G(\mathcal{F}_{t/f})$ we choose from the following options.

- If $F(\mathcal{F}_{t/f}) = \emptyset$ then Definition 2.1.4 requires that $G(\mathcal{F}_{t/f}) = \emptyset$.
- If $\mathcal{F}_{t/f}$ satisfies $F(\mathcal{F}_{t/f}) < \mathcal{F}_{t/f}$ then the permutation $G(\mathcal{F}_{t/f})$ was already determined when $G(F(\mathcal{F}_{t/f}))$ was selected as described above.
- Otherwise assume that $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$. From Definition 2.1.4 the permutation $G(\mathcal{F}_{t/f})$ must map f to f' . We thus choose from the six permutations in S_4 with this property.

In order to reduce the running time of the algorithm, we can prune this recursion so that many gluing permutation selections whose corresponding triangulations do not satisfy the census

constraints are never constructed at all. In particular, the following methods of pruning are available.

- If we are only interested in orientable triangulations then we can reduce the available choices at each stage of the recursion. As we select the various permutations for G we keep track of the orientations imposed upon our tetrahedra. Consider then the situation described above in which we must choose between six available permutations for $G(\mathcal{F}_{t/f})$.

If $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$ for some $f' > 0$ then from the structure of Algorithm 2.2.2 we see that there is some earlier face $\mathcal{F}_{t_0/f_0} < \mathcal{F}_{t/f}$ for which $F(\mathcal{F}_{t_0/f_0}) = \mathcal{F}_{t'/f_0}$. Since we have already selected the permutation $G(\mathcal{F}_{t_0/f_0})$ we have thereby already established an orientation for tetrahedron t' , and furthermore only three of the six permutations available for $G(\mathcal{F}_{t/f})$ will preserve this orientation. Thus we select $G(\mathcal{F}_{t/f})$ from these three permutations only, halving our available choices.

- If we are only interested in prime minimal triangulations then at each stage when a permutation $G(\mathcal{F}_{t/f})$ is selected we ensure that we have not broken any of the constraints listed in Section 2.3.1. If we have broken any of these constraints then we abandon the current choice for $G(\mathcal{F}_{t/f})$ and immediately move on to the next possibility.

Finally, for each gluing permutation selection G that is constructed, we must determine whether G is a minimal representative of its F -isomorphism class. This involves determining whether or not there is some relabelling ρ for which $\rho(G) < G$ and $\rho(F) = F$.

To examine all $(4!)^n n!$ possible relabellings of order n would be infeasibly slow. Recall however that from the results of Algorithm 2.2.2 we already have a list of the automorphisms of F , i.e., the relabellings ρ for which $\rho(F) = F$.

We thus restrict our search to this much smaller list of automorphisms of F . If an automorphism ρ is found for which $\rho(G) < G$ then we discard G . Otherwise G is indeed a minimal representative of its F -isomorphism class and we include it in our final list of results.

2.3.1 Pruning Gluing Permutations

As described above, it is desirable to prune the recursion in Algorithm 2.3.2 so that we avoid generating many gluing permutation selections whose corresponding triangulations do not satisfy our census constraints. Clearly this pruning is dependent upon the specific constraints of the census at hand.

For a census all of whose triangulations must be prime and minimal, the following results regarding edges of low degree can assist in this pruning. Specifically, whenever we select a new permutation $G(\mathcal{F}_{t/f})$ in Algorithm 2.3.2, we compare the resulting partially completed triangulation against the following results. If we have constructed an edge that will force the resulting triangulation to be non-minimal or non-prime then we abandon the current choice for $G(\mathcal{F}_{t/f})$ and try a different permutation instead.

Introducing pruning in a census algorithm to eliminate edges of low degree is a well-known technique. The hyperbolic censuses of Callahan, Hildebrand and Weeks [5, 12] use results regarding low degree edges in hyperbolic triangulations, and the closed orientable censuses of Matveev [29] and Martelli and Petronio [24] use similar results for orientable triangulations.

Theorem 2.3.3 (Degree Three Edges) No minimal triangulation has a non-boundary edge of degree three that belongs to three distinct tetrahedra.

Proof If a triangulation contains an edge of degree three belonging to three distinct tetrahedra, a 3-2 Pachner move as described in Theorem 2.4.1 can be applied to that edge, resulting in a triangulation of the same underlying 3-manifold with fewer tetrahedra. ■

Lemma 2.3.4 No minimal triangulation of a closed prime 3-manifold with ≥ 3 tetrahedra contains an embedded non-separating 2-sphere.

Proof It is a well-known result that any 3-manifold containing an embedded non-separating 2-sphere can be written as a connected sum involving a 2-sphere bundle over the circle. Since our 3-manifold is prime, it follows that the entire 3-manifold must be this 2-sphere bundle, i.e., it must be either the orientable $S^2 \times S^1$ or the non-orientable $S^2 \tilde{\times} S^1$. Both of these 2-sphere bundles however can be triangulated using only two tetrahedra as seen in Sections 3.4.2 and 3.5.1. Thus our triangulation cannot be minimal. ■

Lemma 2.3.5 If a minimal triangulation of a closed prime 3-manifold with ≥ 3 tetrahedra contains an embedded projective plane, then this projective plane must be two-sided and the 3-manifold must be non-orientable.

Proof If the projective plane is one-sided, then a regular neighbourhood of this projective plane has 2-sphere boundary and is in fact $\mathbb{R}P^3$ with a ball removed. Thus our 3-manifold can be written as the connected sum $M \# \mathbb{R}P^3$ for some M . Again primality implies that M is trivial and so our entire 3-manifold is $\mathbb{R}P^3$. But $\mathbb{R}P^3$ can be triangulated using two tetrahedra as seen in Section 3.4.2, and so our triangulation must be non-minimal.

Thus the embedded projective plane is two-sided and so a regular neighbourhood of the projective plane will be non-orientable. Therefore the entire 3-manifold is non-orientable and our result is established. ■

The following result is proven for orientable triangulations by Jaco and Rubinstein in [17]. An extended proof is offered here that also covers non-orientable triangulations.

Theorem 2.3.6 (Degree Two Edges) If a minimal triangulation of a closed prime 3-manifold with ≥ 3 tetrahedra has an edge of degree two, then this 3-manifold is non-orientable with an embedded two-sided projective plane.

Proof Let e be an edge of degree two in a closed prime minimal triangulation with ≥ 3 tetrahedra. If e belongs to only one tetrahedron then it must appear as two distinct edges of that tetrahedron. Figure 2.1 lists the three possible arrangements in which this is possible.

In case I, edge e lies in all four faces of the tetrahedron. However, since e has degree two it can only belong to two faces of the overall triangulation. Thus these four faces are identified in pairs, and the triangulation cannot have more than one tetrahedron.




Figure 2.1: An edge of degree two belonging to only one tetrahedron




Figure 2.2: An edge of degree two belonging to two tetrahedra

In cases II and III, the bottom face of the tetrahedron contains edge e twice and must therefore be identified with some other face of the tetrahedron also containing edge e twice. There is however no other such face and so neither of these cases can occur.

Thus e must belong to two distinct tetrahedra as depicted in Figure 2.2. Consider edges g and h as marked in the diagram. If these edges are identified, the disc between them forms either a 2-sphere or a projective plane. If this disc forms a projective plane then the projective plane is embedded and so Lemma 2.3.5 guarantees the required result.

If the disc between edges g and h forms a 2-sphere, we claim that this 2-sphere must be separating. If the 2-sphere does not intersect itself then Lemma 2.3.4 implies directly that it is separating. Otherwise the only way in which the 2-sphere can intersect itself is if the two vertices at the endpoints of g and h are identified. In this case we deform the 2-sphere slightly by pulling it away from these vertices in the 3-manifold, resulting in an embedded 2-sphere which by Lemma 2.3.4 is again separating.

So if this disc forms a separating 2-sphere then cutting along it splits the underlying 3-manifold into a connected sum decomposition. Assuming the underlying 3-manifold is prime, one side of this disc must thus bound a ball; without loss of generality let it be the side containing faces A and C . The triangulation can then be simplified without changing the underlying 3-manifold by removing this ball and crushing the disc to a single edge, converting Figure 2.2 to a triangular pillow bounded by faces B and D .

If faces B and D are identified, the 3-manifold must be either the 3-sphere or $L(3, 1)$ (the two spaces obtainable by identifying the faces of a triangular pillow) and our triangulation must be non-minimal since each of these spaces can be realised with ≤ 2 tetrahedra as seen in Sections 3.4.1 and 3.4.2. If faces B and D are not identified, the entire pillow can be crushed to a face and our triangulation has been reduced by two tetrahedra and was thus non-minimal.

The only case remaining is that in which edges g and h are not identified. In this case we may crush the disc between edges g and h to an edge, converting Figure 2.2 to a pair of triangular pillows, one bounded by faces A and C and the other bounded by faces B and D .

Each of these pillows may now be crushed to a face in turn. The underlying 3-manifold is only ever changed if we attempt to crush a pillow whose top face is identified with the bottom,

in which case our 3-manifold must again be S^3 or $L(3,1)$ and so our triangulation must be non-minimal. Otherwise we have reduced our triangulation by two tetrahedra and once more it must be non-minimal. ■

The following result is again proven for orientable triangulations by Jaco and Rubinstein in [17] using normal surfaces. A different proof is offered here that uses simpler techniques and once more offers insight into the non-orientable case.

Theorem 2.3.7 (Degree One Edges) If a minimal triangulation of a closed prime 3-manifold with ≥ 3 tetrahedra has an edge of degree one, then this 3-manifold is non-orientable with an embedded two-sided projective plane.

Proof The only way to create an edge of degree one in a 3-manifold triangulation is to snap two faces of a tetrahedron together around the edge between them, as illustrated in Figure 2.3 (where e is the edge of degree one).




Figure 2.3: An edge of degree one

Since our triangulation has ≥ 3 tetrahedra, the two remaining faces of this tetrahedron cannot be identified with each other. So let the upper face be joined to some other tetrahedron, as illustrated in the left hand diagram of Figure 2.4 (note that for simplicity edge e is no longer drawn).




Figure 2.4: An edge of degree one and an adjoining tetrahedron

Again consider edges g and h . If these edges are identified, the disc between them forms either a 2-sphere or a projective plane. As with the proof of Theorem 2.3.6, if it forms a projective plane then this projective plane is embedded and Lemma 2.3.5 provides us with our required result.

If the disc between g and h forms a 2-sphere (which can again be made embedded by pulling it away from the vertices in the underlying 3-manifold), Lemma 2.3.4 implies that this 2-sphere is separating. Cutting along the disc between g and h therefore splits our 3-manifold into two separate pieces each with 2-sphere boundary. Flattening this disc to a single edge effectively fills

in each of these 2-sphere boundary components with balls, producing closed 3-manifolds M and N for which our original 3-manifold is the connected sum $M \# N$.

This procedure is illustrated in the middle diagram of Figure 2.4, where M includes the portion above edge AC (including vertex B) and N includes the portion below edge AC (including vertex D). Note that the portion of the diagram between vertices A , C and B is now a triangular pillow. Furthermore, the portion between vertices A , C and D can be retriangulated using a single tetrahedron with two faces snapped together, much like our original Figure 2.3.

Since our original 3-manifold is prime but is also expressible as the connected sum $M \# N$, it follows that either M or N is in fact this original 3-manifold. We can thus throw away the other of these two components with the aim of forming a new triangulation of our original 3-manifold that uses fewer tetrahedra.

If M is our original 3-manifold, we begin by flattening the triangular pillow ABC to a single face as illustrated in the right hand diagram of Figure 2.3. The only way in which this can alter the underlying 3-manifold is if the two faces bounding this pillow are identified in the original triangulation. In this case however, M must be either the 3-manifold S^3 or $L(3, 1)$ (the two 3-manifolds obtainable by identifying the two faces of a triangular pillow), both of which can be triangulated using ≤ 2 tetrahedra as seen in Sections 3.4.1 and 3.4.2. Otherwise by flattening the pillow ABC and removing the trivial 3-manifold N we can reduce the triangulation of M to a triangulation of the same 3-manifold that uses strictly fewer tetrahedra. Either way we see that our original triangulation is non-minimal.

Alternatively assume that N is our original 3-manifold. In this case we simply discard M and retriangulate the region between vertices A , C and D using a single tetrahedron as described earlier. Once more we obtain a new triangulation of our original 3-manifold that uses strictly fewer tetrahedra and so we see that our original triangulation is non-minimal.

The only remaining case is that in which edges g and h are not identified at all. Here we may crush the disc between g and h to an edge without altering the underlying 3-manifold, as illustrated in the central diagram of Figure 2.4. The region between vertices A , C and D may then be retriangulated using a single tetrahedron as before.

In this case we once more observe that the region between vertices A , C and B becomes a triangular pillow. Furthermore, as a result of the common edges AB and BC , it is impossible to identify the two faces bounding this pillow under any rotation or reflection without identifying edges g and h as a result. Hence the two faces bounding this pillow are distinct and we may flatten the pillow to a face as illustrated in the right-hand diagram of Figure 2.4. This results in a new triangulation of our original 3-manifold with one fewer tetrahedron, and so once more our original triangulation is non-minimal. ■

When considering only orientable triangulations, Theorems 2.3.6 and 2.3.7 can be combined into the following simple result.

Corollary 2.3.8 No minimal triangulation of a prime orientable 3-manifold with ≥ 3 tetrahedra has a non-boundary edge of degree one or two.

2.4 Processing Triangulations

Finally we arrive at step 3 of Algorithm 2.1.16, in which we construct a triangulation from its associated face pairing and gluing permutation selection and determine whether it satisfies the full set of census constraints. Constructing the triangulation is straightforward and is detailed in Theorem 2.1.7. In this section we therefore concentrate on the process of determining whether a given triangulation is suitable for inclusion in the census.

Clearly this process is dependent upon the specific census constraints. The techniques discussed here relate to a census in which only prime minimal triangulations are required.

Note that the methods detailed in Section 2.3.1 eliminate many non-prime and/or non-minimal triangulations from our census before they are ever constructed. We will of course still construct triangulations that are either non-prime or non-minimal; we examine in the following sections a variety of techniques that can identify a large number of these. Section 2.4.4 discusses the ways in which we deal with the remaining triangulations that these techniques do not eliminate.

2.4.1 Elementary Moves

It is often possible to make a local modification to a triangulation that reduces the number of tetrahedra but preserves the underlying 3-manifold, with no knowledge whatsoever of the global triangulation structure or of any properties of the 3-manifold. Such local modifications are referred to as *elementary moves*.

For each of the elementary moves listed below, it is relatively straightforward to test whether the move may be made. In particular, if such a move can be made then the overall triangulation is clearly non-minimal. Note that no properties whatsoever are assumed of the underlying triangulation; these moves may for instance be used with non-orientable triangulations or triangulations with boundary faces or unusual vertex links (such as tori or Klein bottles).

The elementary moves described below are not new, and many of these moves were implemented in 1999 by David Letscher in the program *Normal*, discussed in Section 1.3.1.

Theorem 2.4.1 (3-2 Pachner Move) If a triangulation contains a non-boundary edge of degree three that belongs to three distinct tetrahedra then the following move can be made.

These three tetrahedra (adjacent along our edge of degree three) are replaced with a pair of tetrahedra adjacent along a single face, as illustrated in Figure 2.5. This is called a *3-2 Pachner move* and is described in [34]. This move preserves the underlying 3-manifold and reduces the number of tetrahedra in its triangulation by one.




Figure 2.5: A 3-2 Pachner move

Proof Since the modifications all take place within the boundary of the illustrated polyhedron and since none of the vertex or edge links on the boundary are changed, it is clear that the underlying 3-manifold is preserved. ■

Theorem 2.4.2 (2-0 Vertex Move) Let V be a non-boundary vertex of degree two in a triangulation. Assume the link of V is a 2-sphere and that the two tetrahedra meeting V are distinct. Furthermore, assume that these two tetrahedra meet along three different faces all of which contain V , as illustrated in the left hand diagram of Figure 2.6. Assume that the faces in each tetrahedron opposite V are distinct and are not both boundary faces (though one of them may be a boundary face).




Figure 2.6: A 2-0 vertex move

Then these two tetrahedra can be crushed to a single face, as illustrated in the right hand diagram of Figure 2.6. This is called a *2-0 vertex move*. This move preserves the underlying 3-manifold and reduces the number of tetrahedra in its triangulation by two (and reduces the number of vertices by one).

Proof Let the two faces opposite vertex V be F and G . Since F and G are distinct and not both boundary faces, there is some tetrahedron Δ adjacent to one of these faces that is distinct from the two tetrahedra already seen. Without loss of generality let this new tetrahedron Δ be adjacent to face F as illustrated in the left hand diagram of Figure 2.7.




Figure 2.7: A 2-0 vertex move explained

The result of performing the 2-0 vertex move is illustrated in the right hand diagram of Figure 2.7. The modifications all take place within the boundary of the illustrated polyhedron (the polyhedron bounded on the top by three faces of Δ and on the bottom by G) and none of the vertex or edge links on the boundary of this polyhedron are changed. Thus the 2-0 vertex move preserves the underlying 3-manifold. ■

Theorem 2.4.3 (2-0 Edge Move) Let e be a non-boundary edge of degree two in a triangulation, as illustrated in the left hand diagram of Figure 2.8. Assume that the two tetrahedra meeting e are distinct. Assume that the edges opposite e in each tetrahedron, labelled g and h in the

diagram, are distinct and are not both boundary edges (though one of them may be a boundary edge).




Figure 2.8: A 2-0 edge move

Consider now the four faces in the diagram that do not contain edge e (these are the faces on either side of edges g and h); label these faces G_1 , G_2 , H_1 and H_2 as illustrated in the diagram. Observe that edge g lies between faces G_1 and G_2 and edge h lies between faces H_1 and H_2 . Assume that faces G_1 and H_1 are distinct and that faces G_2 and H_2 are distinct. Assume that all four of these faces are not identified in pairs (though we may have two of these faces identified, such as G_1 and G_2). Assume that we do not simultaneously have two of these faces identified with each other and the other two of these faces as boundary faces.

Then these two tetrahedra can be crushed to a pair of faces, as illustrated in the right hand diagram of Figure 2.8. This is called a *2-0 edge move*. This move preserves the underlying 3-manifold and reduces the number of tetrahedra in its triangulation by two.

Proof Consider the disc bounded by edges g and h that slices through our two tetrahedra. Since g and h are distinct and not both boundary edges, we may crush this disc to a single edge without changing the underlying 3-manifold. After this crushing we are left with two triangular pillows joined along a single edge as illustrated in the left hand diagram of Figure 2.9. We may then retriangulate each of these pillows using two tetrahedra as illustrated in the central diagram of Figure 2.9. As a result of these modifications we have increased our overall number of tetrahedra by two.




Figure 2.9: The intermediate stages of a 2-0 edge move

Our aim is now to collapse each of these pillows to a face using a 2-0 vertex move as described in Theorem 2.4.2. Recall that faces G_1 and H_1 are not identified. Recall also that g and h are not both boundary edges and so G_1 and H_1 cannot both be boundary faces. Thus the pillow

bounded by faces G_1 and H_1 satisfies the conditions of Theorem 2.4.2. We can therefore perform a 2-0 vertex move upon this pillow as illustrated in the right hand diagram of Figure 2.9 without changing the underlying 3-manifold.

Consider now the pillow bounded by faces G_2 and H_2 . If neither of these faces was originally identified with either G_1 or H_1 then a similar argument allows us to perform a 2-0 vertex move upon this second pillow. It is possible however that one of G_2 and H_2 was identified with either G_1 or H_1 , in which case the previous 2-0 vertex move may have changed properties of faces G_2 and H_2 . In this case we must reexamine the conditions of Theorem 2.4.2.

Say then that faces G_2 and H_2 are now identified. Since the conditions of this theorem state that G_2 and H_2 were distinct before we began our modifications, it must be the case that they were identified as a result of the previous 2-0 vertex move that merged G_1 and H_1 into a single face. This implies that both faces G_2 and H_2 were originally identified with G_1 and H_1 in some order, which is impossible since the conditions of this theorem state that faces G_1, G_2, H_1 and H_2 are not identified in pairs.

Finally consider the situation in which faces G_2 and H_2 are both boundary faces. Since g and h were not originally both boundary edges it follows that G_2 and H_2 were not originally both boundary faces. It must then be the case that this situation was created as a result of the previous 2-0 vertex move. Specifically the original scenario must have involved one of G_2 and H_2 being a boundary face, the other of G_2 and H_2 being identified with one of G_1 and H_1 , and the other of G_1 and H_1 also being a boundary face. This scenario however is explicitly excluded by conditions of this theorem.

Thus after our first 2-0 vertex move the pillow bounded by G_2 and H_2 still satisfies the conditions of Theorem 2.4.2, and so we can perform an additional 2-0 vertex move upon this second pillow without changing the underlying 3-manifold. This completes the 2-0 edge move with no changes in the underlying 3-manifold and an overall reduction of two tetrahedra. ■

Theorem 2.4.4 (2-1 Edge Move) Let e be a non-boundary edge of degree one in a triangulation, as illustrated in the left hand diagram of Figure 2.10. Label the single tetrahedron containing e as T , label the endpoints of e as A and B and label the remaining two vertices of T as C and D .




Figure 2.10: A 2-1 edge move

Assume that face CAD is not a boundary face and let S be the tetrahedron adjacent to T along this face. Assume that tetrahedra T and S are distinct and let the vertex of S that does not belong to face CAD be labelled as E . Assume that edges CE and DE of tetrahedron S are distinct (these edges are labelled g and h on the diagram) and that these edges are not both boundary (though one of them may be boundary).

Then the two tetrahedra T and S can be merged into a single tetrahedron, with the region between edges g and h and vertex A flattened to a single face. This operation is called a *2-1 edge move* and is illustrated in the right hand diagram of Figure 2.10. This move preserves the underlying 3-manifold and reduces the number of tetrahedra in its triangulation by one.

Proof Consider faces CAE and DAE . Since g and h are not both boundary edges it follows that CAE and DAE cannot both be boundary faces. Furthermore, assume that faces CAE and DAE are identified. Regardless of which of the six possible rotations or reflections is used for this identification, the presence of edge AE and edge CA (or equivalently DA) in both faces would result in the identification of edges g and h . Because edges g and h are known to be distinct it follows that faces CAE and DAE must also be distinct.

At this point we employ a strategy similar to that used in the proof of Theorem 2.4.3. Consider the disc bounded by edges g and h that slices through both tetrahedra S and T . Since edges g and h are distinct and not both boundary we can crush this disc to a single edge without changing the underlying 3-manifold. This reduces the region between edges g and h and vertex A to a triangular pillow which we retriangulate using two tetrahedra. The region between edges g and h and vertex B is retriangulated using a single tetrahedron with two of its faces identified. The resulting structure is illustrated in the central diagram of Figure 2.11.




Figure 2.11: The intermediate stage of a 2-1 edge move

Using the earlier observation that faces CAE and DAE are distinct and not both boundary, we can invoke Theorem 2.4.2 to perform a 2-0 vertex move upon the triangular pillow, reducing it to a single face. This completes the 2-1 edge move with no changes in the underlying 3-manifold and an overall reduction of one tetrahedron. ■

2.4.2 0-Efficiency

Jaco and Rubinstein develop in [17] a theory of 0-efficient triangulations. 0-efficiency is defined below and presents a variety of implications for the theory of minimal triangulations. Several of the results presented in [17] are outlined in this section.

In the context of constructing a census of triangulations, the most important result of [17] is Corollary 2.4.7 which shows that 0-efficiency can be used as an algorithmic test for concluding that certain triangulations are either non-prime or non-minimal. Details of a method for testing for 0-efficiency are included in Algorithm 2.4.13.

0-efficiency is extremely slow to compute in comparison to the other tests described in Section 2.4 since it requires the calculation of the vertex normal surfaces of a triangulation (although

David Letscher and others have offered methods for substantially increasing the speed of this calculation in various seminars and informal discussions).

Because of this computational difficulty, we additionally present a number of simple consequences that Jaco and Rubinstein draw from their core results. These consequences allow for fast tests that in some cases allow us to conclude that a triangulation is non-prime or non-minimal without having to run a full 0-efficiency test. Of particular interest are Corollary 2.4.9 which tests the number of vertices in a triangulation and Corollary 2.4.12 which tests for faces whose edges are identified to form a cone.

At this point we present a definition of 0-efficiency for a closed 3-manifold triangulation followed by some of the central results of [17]. Although Jaco and Rubinstein also discuss 0-efficiency for triangulations with boundary faces and ideal triangulations, such results are not used here.

Definition 2.4.5 (0-Efficiency) Let T be a triangulation of a closed 3-manifold. T is said to be *0-efficient* if the only embedded normal 2-spheres in T are vertex linking (i.e., contain no quadrilateral discs).

Theorem 2.4.6 Let T be a minimal triangulation of a closed orientable irreducible 3-manifold. Then either T is 0-efficient or T is a triangulation of $\mathbb{R}P^3$ or $L(3, 1)$.

Proof The proof given in [17] essentially runs as follows. If T is not 0-efficient, it contains a non-vertex-linking embedded normal 2-sphere. We crush this 2-sphere to a point in T , an operation which reduces the number of tetrahedra in T and has the side-effect of decomposing T along this 2-sphere. Since our 3-manifold is irreducible, this decomposition must have simply split off a 3-sphere which we can throw away. We now have a smaller triangulation of the same 3-manifold, and we repeat this process until our triangulation has become 0-efficient.

Of course there are many more details and special cases to consider in this proof. The exceptions $\mathbb{R}P^3$ and $L(3, 1)$ arise from these special cases. ■

Corollary 2.4.7 Every minimal triangulation of a closed orientable prime 3-manifold with ≥ 3 tetrahedra is 0-efficient.

Proof This follows immediately from Theorem 2.4.6 when we observe that the only closed orientable irreducible 3-manifold that is not prime is $S^2 \times S^1$, and that $S^2 \times S^1$, $\mathbb{R}P^3$ and $L(3, 1)$ each have triangulations with ≤ 2 tetrahedra (as seen in Sections 3.4.1 and 3.4.2). ■

Corollary 2.4.7 is a strong result and many of the other minimality results discussed in this chapter can (with a little manipulation) be seen merely as special cases of this theorem.

Following the presentation of their core results, Jaco and Rubinstein prove a variety of properties of 0-efficient triangulations. When combined with Theorem 2.4.6 these properties of 0-efficient triangulations become properties of minimal triangulations, some of which are quite simple to test.

Theorem 2.4.8 Let T be a 0-efficient triangulation of a closed orientable 3-manifold. Then either T has only one vertex or T is a two-vertex triangulation of S^3 .

Proof The proof in [17] involves finding an edge in T whose endpoints are distinct vertices; if such an edge exists then let S be the 2-sphere that bounds a small neighbourhood of this edge. It is proven that either S can be manipulated into a non-vertex-linking embedded normal 2-sphere, or that S bounds a ball on the side away from our edge and hence T is a triangulation of S^3 (since S also bounds a ball on the side towards our edge).

In the case where T triangulates S^3 with more than two vertices, the arguments are more complex but again involve constructing a 2-sphere that can be manipulated into a non-vertex-linking embedded normal 2-sphere. This is done by locating barriers in the 1-skeleton of T that prevent this 2-sphere from being reduced all the way down to one or more vertex links. ■

Corollary 2.4.9 Every minimal triangulation of a closed orientable prime 3-manifold with ≥ 3 tetrahedra has only one vertex.

Proof This follows immediately from Corollary 2.4.7 and Theorem 2.4.8, making the observation that the special case S^3 has a triangulation with only one tetrahedron (see Section 3.4.1). ■

Lemma 2.4.10 Let T be a 0-efficient triangulation of a closed orientable 3-manifold that is not S^3 . Then no single edge of T bounds an embedded disc in the underlying 3-manifold.

Proof The proof of this result in [17] is similar to that of Theorem 2.4.8. This time, if such a disc exists, we let S be the 2-sphere bounding a small neighbourhood of this disc, and again we prove that either S can be converted into a non-vertex-linking embedded normal 2-sphere or that S bounds a ball on both sides. ■

Recall Theorem 2.3.7 which concerns minimal triangulations and degree one edges. When restricted to orientable triangulations this theorem becomes an immediate corollary of Lemma 2.4.10, since when a tetrahedron is folded onto itself to form a degree one edge, the opposite edge of this tetrahedron bounds an embedded disc slicing through the tetrahedron.

From Lemma 2.4.10 Jaco and Rubinstein prove the following result regarding faces in a triangulation.

Theorem 2.4.11 Let T be a 0-efficient triangulation of a closed orientable 3-manifold that is not S^3 . Then no face of T has two of its edges identified to form either a cone or a pinched cone as illustrated in Figure 2.12. In both the cone and the pinched cone note that the third edge of the face is not identified with the others.




Figure 2.12: A face with two edges identified to form a cone and a pinched cone

Proof If the two vertices of the cone are distinct, the cone itself is an embedded disc bounded by the edge at its base. If the vertices of the cone are identified (i.e., the cone is pinched), it is

described in [17] how this cone may be manipulated slightly at its tip to remove the self-intersection so that Lemma 2.4.10 can again be applied. ■

Corollary 2.4.12 No minimal triangulation of a closed orientable prime 3-manifold with ≥ 3 tetrahedra has a face in which two edges are identified to form either a cone or a pinched cone (again with the third edge of the face not identified with the others).

Proof Once more we use Corollary 2.4.7 and Theorem 2.4.11, again observing that S^3 can be triangulated with only one tetrahedron (Section 3.4.1). ■

We close this section with a proof that 0-efficiency can be tested simply by evaluating vertex normal surfaces.

Algorithm 2.4.13 (0-Efficiency Test) Let T be a triangulation of a closed orientable 3-manifold and let \mathcal{V} be the set of vertex embedded normal surfaces in T . Then T is 0-efficient if and only if \mathcal{V} contains no one-sided projective planes and \mathcal{V} contains no 2-spheres with one or more quadrilateral discs.

Proof Let T be 0-efficient. Then the only embedded normal 2-spheres in T are vertex linking spheres which contain only triangular discs; thus \mathcal{V} contains no 2-spheres with quadrilateral discs. Furthermore, say \mathcal{V} contains a one-sided projective plane p . Then the normal surface $2p$ (the surface whose vector is double that of p) is the boundary of a regular neighbourhood of p , i.e., $2p$ is a 2-sphere. Since p is not vertex linking (as all vertex links are spheres), $2p$ is certainly not vertex linking and we see that T could not have been 0-efficient.

Suppose therefore that T is not 0-efficient. Then there is some non-vertex-linking embedded normal 2-sphere v in T . We can express v as the combination $\lambda_1 u_1 + \dots + \lambda_k u_k$ where each u_i is a vertex embedded normal surface in T and each $\lambda_i > 0$. However, since Euler characteristic is a linear function of the normal surface vector, this implies that $\chi(v) = \lambda_1 \chi(u_1) + \dots + \lambda_k \chi(u_k)$ and in particular since $\chi(v) > 0$ we must have $\chi(u_i) > 0$ for some i . Thus (since all vertex normal surfaces are connected) there is some vertex embedded normal surface u_i that is either a 2-sphere or a projective plane.

If this u_i is a 2-sphere and $k > 1$, u_i cannot be vertex linking since the sum of a vertex link with any other normal surface is disconnected. If this u_i is a 2-sphere and $k = 1$ then $u_i = v$ and so again u_i is not vertex linking. Either way we have a non-vertex-linking embedded 2-sphere in \mathcal{V} which must in turn have quadrilateral discs (since the only connected normal surfaces with no quadrilateral discs are the vertex links).

Otherwise this u_i must be a projective plane. In this case it must be a one-sided projective plane since otherwise the underlying 3-manifold would be non-orientable. Thus we have a one-sided embedded projective plane in \mathcal{V} . ■

2.4.3 Special Subcomplexes

Here we identify a variety of simple structures whose presence within a triangulation can indicate that the triangulation is either non-prime or non-minimal. In particular the structures described in this section are formed from a small number of faces in a triangulation. Searching for these

structures is quite fast and can allow us to exclude certain triangulations from a census without having to run more expensive tests such the 0-efficiency test described in Section 2.4.2.

The specific results that lead to simple algorithmic tests are Theorem 2.4.14 which involves searching for faces with all three edges identified and Corollaries 2.4.17 and 2.4.18 which involve searching for 2-spheres formed from a pair of faces.

One can prove that in many cases a triangulation that contains one of these structures will also contain a non-trivial normal 2-sphere and thus fail the more general 0-efficiency test described in Corollary 2.4.7. However, in addition to the speed advantage mentioned above, these specialised subcomplex tests are of interest because some of them do not rely on orientability and can thus be used in more general situations.

Theorem 2.4.14 If a closed orientable triangulation with ≥ 3 tetrahedra contains a face whose three edges are all identified with each other in the same direction around the face, as illustrated in Figure 2.13, then the triangulation cannot be both prime and minimal.




Figure 2.13: A face with three identified edges

Proof Consider a regular neighbourhood of the face in question; this will be a triangular prism with the original face slicing through the centre of the prism and identifications between the half-rectangular faces on either side. In order to maintain orientability of the 3-manifold and preserve the edge identifications of the original face, there is only one way in which these half-rectangles can be identified (excluding rotations and reflections of the prism); this is illustrated in Figure 2.14. In particular, faces $XYBA$ and $ZXDF$ are identified, faces $YZCB$ and $XYED$ are identified and faces $ZXAC$ and $YZFE$ are identified, with XYZ as the original face.




Figure 2.14: A regular neighbourhood of a face with three identified edges

It can be seen then that the boundary of this prism is a 2-sphere and the body of the prism forms a lens space $L(3,1)$ with a removed ball. Thus our original triangulation represents the 3-manifold $L(3,1)\#M$ for some M .

Hence our triangulation is either non-prime or it is simply $L(3,1)$. Since there is a triangulation of $L(3,1)$ with two tetrahedra (see Section 3.4.2) and our triangulation has ≥ 3 tetrahedra, it follows that our triangulation is either non-prime or non-minimal. ■

Theorem 2.4.15 Let T be a closed triangulation with ≥ 3 tetrahedra. Consider two faces F_1 and F_2 of T that are joined along at least one edge.

Slicing T along F_1 and F_2 will produce a new (possibly disconnected) triangulation with four boundary faces; call this T' . Note that T' might have vertices whose links are neither spheres nor discs, and so might not actually represent one or more 3-manifolds.

If T' contains multiple boundary components (as opposed to a single four-face boundary component) and one of these boundary components is a sphere as illustrated in Figure 2.15, where the two faces of the sphere correspond to faces F_1 and F_2 in the original triangulation T , then T is either non-prime or non-minimal.




Figure 2.15: A sphere formed by identifying the boundaries of two triangles

Note that it does not matter if some vertices of this boundary sphere are identified (thus producing a sphere pinched at two or more points) – this is indeed expected if T' has non-standard vertex links as suggested above. Only the edge identifications illustrated in Figure 2.15 are required for the conditions of this theorem to be met.

Proof Let T be a prime minimal triangulation satisfying all of the conditions described in the theorem statement. We first observe that any embedded 2-sphere contained within T must bound a ball. Let S be some embedded 2-sphere within T . From Lemma 2.3.4 we see that S must be a separating 2-sphere which thus breaks T into a connected sum. Since T is prime it follows that one of the terms of this connected sum must be trivial and so S bounds a ball on one side.

Consider now the triangulation T' as described in the theorem statement. Since T' has four boundary faces and multiple boundary components, it must have precisely two boundary components each with two faces (since a boundary component must have an even number of faces). Let these boundary components be ∂_1 and ∂_2 where ∂_1 is a sphere as illustrated in Figure 2.15. Let R_1 and R_2 represent the regions of T' just inside boundary components ∂_1 and ∂_2 respectively. Finally let S be an embedded sphere located in R_1 just behind the spherical boundary component ∂_1 . This scenario is illustrated in Figure 2.16.




Figure 2.16: Regions within the split triangulation T'

As noted in the theorem statement, vertices of the boundary sphere ∂_1 might be identified and thus it may in fact be impossible to place the sphere S entirely within region R_1 . If this is the

case then we simply push S outside region R_1 in a small neighbourhood of each offending vertex. Examples of the resulting sphere S are illustrated for two different cases in Figure 2.17.




Figure 2.17: Pushing sphere S outside region R_1 in the neighbourhood of a vertex

We can now view S as an embedded 2-sphere in the original triangulation T . Note that any situation in which S is pushed outside region R_1 as described above simply results in S being pushed into a small neighbourhood of a vertex in region R_2 . In particular, since the link of each vertex in T is a 2-sphere, such an operation does not induce any self-intersections in S .

From our earlier observations we see that the embedded 2-sphere S must bound a ball in T . We take cases according to whether this ball lies on the side of S including region R_1 or region R_2 .

- Suppose that S bounds a ball on the side containing region R_2 . In this case we can remove the component of T' containing region R_2 and replace it with a two-tetrahedron triangular pillow as illustrated in Figure 2.18. This pillow can then be joined to region R_1 along the spherical boundary ∂_1 resulting in a new triangulation T'' whose underlying 3-manifold is identical to that of T .




Figure 2.18: A replacement two-tetrahedron triangular pillow

- Otherwise S bounds a ball on the side containing region R_1 (excluding a small neighbourhood of the boundary ∂_1). In this case we remove the component of T' containing region R_1 and again replace it with the two-tetrahedron triangular pillow illustrated in Figure 2.18 which has a spherical boundary identical to ∂_1 . Once more we obtain a new triangulation T'' whose underlying 3-manifold is identical to that of T .

Note that even if the component of T' containing region R_1 has a vertex whose link is a multiply-punctured sphere, the underlying 3-manifolds of T and T'' remain identical. Such a scenario is illustrated in the left hand diagram of Figure 2.19, with the corresponding triangular pillow illustrated in the right hand diagram of the same figure. Since S is pushed away from the vertex in question, it can be seen that in both triangulations T and T'' sphere S bounds a ball on the side containing region R_1 . Likewise, S bounds the same 3-manifold on the side containing region R_2 in both triangulations.

We see then that in each of the above cases we have removed at least one tetrahedron and inserted two, resulting in a net increase of at most one tetrahedron. However, since the new triangulation T'' contains no boundary faces and the two original faces F_1 and F_2 were distinct,




Figure 2.19: A case involving a vertex whose link is a multiply-punctured sphere

we can perform a 2-0 vertex move upon this triangulation as described in Theorem 2.4.2. The resulting triangulation has the same underlying 3-manifold as T and T'' but now uses at least one fewer tetrahedron than the original triangulation T . Thus triangulation T cannot be minimal. ■

Theorem 2.4.16 Let T be a closed orientable triangulation with ≥ 3 tetrahedra. As in Theorem 2.4.15, consider two faces F_1 and F_2 of T that are joined along at least one edge.

Again slicing T along F_1 and F_2 will produce a new (possibly disconnected) triangulation T' with four boundary faces whose vertices might have links that are neither spheres nor discs.

If T' contains multiple boundary components (as opposed to a single four-face boundary component) and one of these boundary components is a sphere as illustrated in Figure 2.20, where the two faces of the sphere correspond to faces F_1 and F_2 in the original triangulation T , then T is either non-prime or non-minimal.




Figure 2.20: A sphere formed by creating two cones and identifying their boundaries

Again it does not matter if some vertices of this boundary sphere are identified, thus producing a sphere pinched at two or more points. Only the edge identifications illustrated in Figure 2.20 are required.

Proof The proof of this result is almost identical to the proof of Theorem 2.4.15. The only difference lies in the retriangulation of whichever component of T' forms a ball. Instead of using the triangular pillow illustrated in Figure 2.18, we now use a ball formed from a single tetrahedron with two of its faces identified as illustrated in Figure 2.21.




Figure 2.21: A replacement one-tetrahedron ball

The new triangulation T'' as described in the proof of Theorem 2.4.15 is now constructed by removing at least one tetrahedron and inserting only one, and so T'' represents the same 3-manifold as T and uses at most the same number of tetrahedra.

If T'' uses strictly fewer tetrahedra than T then triangulation T cannot be minimal. Otherwise T'' contains at least three tetrahedra, and furthermore since it includes the ball illustrated in

Figure 2.21 it also contains an edge of degree one. Since T'' is a closed orientable prime triangulation we can conclude from Theorem 2.3.7 that T'' is non-minimal, and since T contains the same number of tetrahedra as T'' it follows that the original triangulation T is likewise non-minimal. ■

Corollary 2.4.17 Let T be a closed triangulation with ≥ 3 tetrahedra. If T contains two faces whose edges are identified to form a sphere as illustrated in Figure 2.15 and the three edges of this sphere are distinct (i.e., none are identified with each other in the triangulation), then T is either non-prime or non-minimal.

Proof Since the three edges of this sphere are distinct, the sphere has no self-intersections except possibly at the vertices of its two faces. Combined with the fact that all spheres in 3-manifolds are two-sided, we see that slicing T along the two faces described produces two spherical boundary components that are also triangulated as illustrated in Figure 2.15. Thus Theorem 2.4.15 can be invoked and T must be non-prime or non-minimal. ■

Corollary 2.4.18 Let T be a closed orientable triangulation with ≥ 3 tetrahedra. If T contains two faces whose edges are identified to form a sphere as illustrated in Figure 2.20 and the three edges of this sphere are distinct (i.e., none are identified with each other in the triangulation), then T is either non-prime or non-minimal.

Proof As in Corollary 2.4.17, the distinctness of the three edges of the sphere implies that if we slice T along the two faces described we obtain two boundary components each triangulated according to Figure 2.20. The conditions of Theorem 2.4.16 are thus satisfied and we obtain the desired result. ■

2.4.4 Remaining Triangulations

Although the techniques of Sections 2.4.1, 2.4.2 and 2.4.3 can be used to exclude a variety of triangulations, we have not yet seen a criterion that allows us to determine that a particular triangulation can for certain be included in a list of census results. We thus turn to this final stage of processing in the census algorithm.

It is assumed that we are constructing a census of prime minimal triangulations whose underlying 3-manifolds satisfy some additional set of constraints. It is also assumed that the results of each census with similar constraints but a smaller number of tetrahedra are available.

Consider then a triangulation T that has been constructed using the first two steps of Algorithm 2.1.16 and has not been identified as non-prime or non-minimal by any of our earlier techniques. We must make a final decision as to whether T will or will not be included in the census.

The following list outlines the different methods that are available for arriving at such a decision. Note that some of these methods can be fully or partially automated and some require human interaction and/or intuition.

- **Simplification:** We can find a way of simplifying the triangulation so that it represents the same underlying 3-manifold but is constructed from fewer tetrahedra. In this case the triangulation is non-minimal and we exclude it from our census.

Simplification is attempted using elementary moves such as those discussed in Section 2.4.1. However, instead of simply testing for a single move that reduces the number of tetrahedra as described in Section 2.4.1, we also experiment with longer sequences of moves including moves that preserve or temporarily increase the number of tetrahedra. This increases our chances of simplifying the triangulation, since moves that temporarily increase the number of tetrahedra are sometimes necessary in order for a triangulation to be ultimately simplified.

Additional moves that are used but not discussed in Section 2.4.1 include the 2-3 Pachner move (the inverse of the 3-2 Pachner move described in Theorem 2.4.1) and the 4-4 move which involves replacing four distinct tetrahedra about an edge of degree four with a differently arranged set of four tetrahedra.

It is worth noting that David Letscher implemented random searches for simplifying sequences of moves in *Normal*, the precursor to *Regina* discussed in Section 1.3.1.

- **Invariant Analysis:** We can evaluate a variety of 3-manifold invariants of our triangulation, such as homology and homotopy groups. If no 3-manifold with the same values for these invariants has appeared in a census involving fewer tetrahedra and if our triangulation satisfies all of the census constraints with the exception of minimality, then our triangulation is indeed minimal and should be included in our census.

Of the remaining census constraints, primality is in some cases particularly difficult to test for. In the orientable case however ensuring primality is straightforward since it is proven by Kneser [21] and Schubert [42] that any triangulation of a reducible 3-manifold contains an embedded essential normal 2-sphere. Thus all non-prime orientable 3-manifolds are eliminated by the 0-efficiency test described in Section 2.4.2.

It is worth noting that Matveev in [30] reports remarkable success with the invariants of Turaev and Viro introduced in [46]. Calculation of these invariants is a planned extension for the program *Regina*.

Note that even if a triangulation is marked for inclusion in our census using invariant analysis, the underlying 3-manifold must still be identified for the census results to be useful. This identification process is discussed below.

- **Identification:** We can identify the underlying 3-manifold of our triangulation. If the 3-manifold has already appeared in a census involving fewer tetrahedra or the 3-manifold is not prime then we discard the triangulation. Otherwise the triangulation is prime and minimal, and assuming it satisfies the remaining census constraints we include it in our census.

Identification of the underlying 3-manifold of an arbitrary triangulation is a central problem in low-dimensional topology and currently no algorithm for this task exists. Fortunately our census triangulations generally involve small numbers of tetrahedra and so human analysis and insight in many cases will suffice.

The calculation of 3-manifold invariants can offer strong hints as to what the underlying 3-manifold might be although it cannot prove any such claims. The final claim that triangulation T has underlying 3-manifold M must be proven by showing that T realises a known construction for M . For instance, when identifying a lens space we can search for the

corresponding genus one Heegaard splitting within the triangulation, and when identifying a more general Seifert fibred space we can search for the corresponding fibres within the triangulation.

Although this process of identification is tedious, it can fortunately be done for large classes of 3-manifolds simultaneously. In Section 3.3 we examine several infinite families of triangulations all of whose underlying 3-manifolds are identified. The first step in the identification process is then to test whether our triangulation belongs to one of these infinite families; if it does then the underlying 3-manifold can be identified. This process can be automated and is indeed implemented in *Regina*.

There will in general be triangulations in our census that do not belong to any of the families described in Section 3.3. In this case, once the triangulations have been identified using other means, we search for common properties of these triangulations in the hope of creating new infinite families of recognisable triangulations that can be used in future census runs with larger numbers of tetrahedra.

2.5 Census Statistics

Section 2.4 concludes the full description of the census algorithm as used by the program *Regina* to create the census data discussed in Chapter 3. Improvements to this algorithm have already been designed and in some cases implemented; these are discussed in Sections 2.6 and 2.7. At this point however we pause to examine the efficiency of the census algorithm and in particular the effectiveness of the pruning and processing techniques of Sections 2.3.1 and 2.4. Note that the results listed below pertain only to the census algorithm as it is presented in the previous sections of this chapter, and in particular without the improvements discussed in Sections 2.6 and 2.7.

A variety of statistics for census runs with different constraints and different numbers of tetrahedra are presented in Tables 2.1 and 2.2. Each row of each table describes a census of closed prime minimal triangulations with a particular number of tetrahedra, where only orientable triangulations are considered in Table 2.1 and only non-orientable triangulations are considered in Table 2.2. The results listed were obtained using the program *Regina* which is discussed in more detail in Section 1.3. The meanings of the individual columns of each table are as follows.

- *Tetrahedra*: The number of tetrahedra used for this particular census.
- *Full*: The total number of triangulations produced by the census algorithm if we use none of the pruning techniques of Section 2.3.1 and do not eliminate any triangulations as discussed in Section 2.4. This is precisely the number of isomorphism classes of closed triangulations with the appropriate number of tetrahedra that satisfy the appropriate orientability constraint.
- *Pruned*: The number of triangulations produced by the census algorithm if we use the pruning techniques of Section 2.3.1 during the generation of gluing permutation selections and in addition remove triangulations that do not pass some of the faster automated tests described in Section 2.4.

The specific tests from Section 2.4 that are used include the elementary move tests listed in Section 2.4.1 and the one-vertex test described by Corollary 2.4.9. We also eliminate

triangulations for which the computer can find a sequence of elementary moves that may temporarily increase the number of tetrahedra but nevertheless results in an overall reduction in the number of tetrahedra, as discussed in Section 2.4.4.

- *0-Eff./Sub.*: The number of triangulations listed in the *Pruned* column that, in addition to the tests listed above, pass the slower 0-efficiency test described in Corollary 2.4.7 and the subcomplex tests described in Section 2.4.3.

Note that the 0-efficiency and subcomplex tests pertain only to triangulations with ≥ 3 tetrahedra. Furthermore the 0-efficiency test pertains only to orientable triangulations and so this column is simply labelled *Sub.* in Table 2.2.

- *Final*: The actual number of prime minimal triangulations in the final census.
- *Full Time*: The running time taken to create the census described in the *Full* column. Running times are measured on a single 1.2GHz Pentium III processor and are displayed as *h:mm:ss* unless otherwise indicated. Note that a time of 0:00 simply indicates a running time of less than half a second.
- *Pruned Time*: The running time taken to create the census described in the *Pruned* column. Running times are presented in the same format as in the *Full Time* column.

In addition to the censuses described in Chapter 3, the construction of orientable censuses for seven and eight tetrahedra is partially complete at the time of writing. Although much of the human processing described in Section 2.4.4 remains to be done, the automated processing for these censuses has already been carried out. Partial results for these censuses are thus included in Table 2.1.

The orientable censuses with seven and eight tetrahedra and the non-orientable census with six tetrahedra were considered sufficiently large that a full unpruned census was not carried out. For this reason the corresponding entries in the *Full* and *Full Time* columns remain empty.

Tetrahedra	Full	Pruned	0-Eff./Sub.	Final	Full Time	Pruned Time
1	4	4		4	0:00	0:00
2	16	12		10	0:00	0:00
3	76	7	7	7	0:00	0:00
4	532	17	17	15	0:07	0:03
5	4,807	50	50	40	6:21	2:28
6	52,946	168	167	115	7:44:50	2:49:29
7		569	561			10 days
8		2435	2341			27 months

Table 2.1: Statistics for the closed orientable census

We can observe from Tables 2.1 and 2.2 that the tests and pruning techniques used to obtain the figures in the *Pruned* columns are particularly effective for removing unwanted triangulations from a census. The pruning techniques of Section 2.3.1 in particular are responsible for the noticeable improvement in running time between the *Full Time* and *Pruned Time* columns for the orientable censuses. The running time improvements in the non-orientable censuses are less striking because many of the results in Section 2.3.1 pertain only to orientable triangulations.

Tetrahedra	Full	Pruned	Sub.	Final	Full Time	Pruned Time
1	0	0		0	0:00	0:00
2	1	1		1	0:00	0:00
3	5	3	3	2	0:02	0:02
4	45	9	4	0	3:16	3:03
5	377	33	21	0	6:33:44	6:01:53
6		385	120	24		5 weeks

Table 2.2: Statistics for the closed non-orientable census

For orientable triangulations it is apparent that the slower 0-efficiency test eliminates very few triangulations that the faster tests used for the *Pruned* column do not. It is possible however that this is a consequence of the fact that several of the faster tests can for many triangulations be seen as special cases of the 0-efficiency results, and that as the number of tetrahedra increases a growing number of triangulations will be found that are not identified as non-prime or non-minimal until the more general 0-efficiency test is run.

The subcomplex tests of Section 2.4.3 show themselves to be far more effective in the non-orientable case than in the orientable case at eliminating triangulations that have passed the earlier tests used for the *Pruned* column. This is true despite the fact that the only test from Section 2.4.3 that is applicable to non-orientable triangulations is the 2-sphere test described in Corollary 2.4.17. A possible reason for this observation is the fact that many of the tests used for the *Pruned* column are specific to orientable triangulations only.

2.6 Face Pairings Revisited

Recall from the closing remarks of Section 2.2 that the running time spent by the census algorithm in generating face pairings is negligible. Since the remaining steps of the algorithm (generating gluing permutation selections and constructing and verifying the resulting triangulations) can be performed independently for each face pairing, this leads to an obvious parallelisation of Algorithm 2.1.16.

Specifically, we can distribute this census algorithm across a number of machines or processors as follows. In a short initial calculation on a single machine we generate all face pairings. Each machine or processor that is contributing to the task then repeatedly picks up the next face pairing waiting to be processed, generates its corresponding gluing permutation selections and resulting triangulations as described in Algorithm 2.1.16 and stores its results. This procedure continues with the various machines or processors simultaneously processing different face pairings until each face pairing has been dealt with.

Regina offers tools for this type of parallelisation and the larger census runs described in Section 2.5 were indeed performed by a number of high-performance machines simultaneously. As a result of this parallelisation it became apparent that for most face pairings either no triangulations were constructed or they were all weeded out using the elementary moves described in Section 2.4.1.

Because of this observation it made sense to seek a test that would be capable in a large number of cases of quickly concluding that a face pairing could never lead to a triangulation satisfying the census constraints. Since the face pairing generation itself is so fast, if we could quickly eliminate

$p\%$ of all face pairings in this way then we could thus eliminate approximately $p\%$ of the total census running time.

In this section we develop fast methods of eliminating face pairings that can never lead to closed orientable prime minimal triangulations. As mentioned above, these techniques were discovered as a result of the census runs and so were not used in any of the census runs whose statistics are outlined in Section 2.5 and whose results are discussed in Chapter 3. They will however be incorporated into future census runs as described in Section 2.7.

It is worth noting that Martelli and Petronio prove in [24] different results that likewise restrict the associated face pairings of minimal triangulations with particular properties, though their results are phrased in terms of skeleta of special spines of 3-manifolds. The results presented here were developed independently and neither set of results implies the other.

2.6.1 Face Pairing Graphs

We begin with a visual representation of face pairings as multigraphs, as initially suggested in the remarks following Definition 2.1.2. Recall that a multigraph is a graph in which an edge may join a vertex to itself, and in which multiple edges may join the same pair of vertices.

Definition 2.6.1 Let F be a face pairing of order n . The *face pairing graph* for F is the multigraph on n vertices constructed as follows. We label the vertices of the graph $0, 1, \dots, n-1$, and then for each unordered pair of faces $\mathcal{F}_{t/f}$ and $\mathcal{F}_{t'/f'}$ for which $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$ and $F(\mathcal{F}_{t'/f'}) = \mathcal{F}_{t/f}$ (recalling from Definition 2.1.2 that these two relations are equivalent) we add an edge between vertices t and t' .

Example 2.6.2 Consider the face pairing F of order 2 defined as follows.

$$\begin{array}{llll} F(\mathcal{F}_{0/0}) = \mathcal{F}_{0/1} & F(\mathcal{F}_{0/1}) = \mathcal{F}_{0/0} & F(\mathcal{F}_{0/2}) = \mathcal{F}_{1/0} & F(\mathcal{F}_{0/3}) = \mathcal{F}_{1/1} \\ F(\mathcal{F}_{1/0}) = \mathcal{F}_{0/2} & F(\mathcal{F}_{1/1}) = \mathcal{F}_{0/3} & F(\mathcal{F}_{1/2}) = \emptyset & F(\mathcal{F}_{1/3}) = \emptyset \end{array}$$

The face pairing graph for F is illustrated in Figure 2.22.




Figure 2.22: A sample face pairing graph

We see then that a face pairing graph contains slightly less information than its underlying face pairing. Specifically it details which tetrahedra in a triangulation are joined to which others and along how many faces. It does not however specify which specific faces of each tetrahedron are involved in each such joining.

Fortunately this loss of information does not matter when we restrict our attention to isomorphism classes, as evidenced by the following result.

Lemma 2.6.3 Two face pairings are isomorphic if and only if their face pairing graphs are isomorphic.

Proof Consider face pairings F_1 and F_2 of order n . If these face pairings are isomorphic then $F_2 = \rho(F_1)$ for some relabelling ρ . We recall from Definitions 2.1.10 that a relabelling induces a permutation on the tetrahedra $\{\mathcal{T}_0, \dots, \mathcal{T}_{n-1}\}$, and so we can define a permutation γ on the graph vertices $\{0, \dots, n-1\}$ for which $\gamma(v) = w$ if and only if $\rho(\mathcal{T}_v) = \mathcal{T}_w$.

Let r and s be any two (possibly identical) graph vertices. From Definition 2.6.1 the number of edges between vertices r and s in the face pairing graph for F_1 is the number of unordered pairs of faces $\mathcal{F}_{r/f}$ and $\mathcal{F}_{s/f'}$ for which $F_1(\mathcal{F}_{r/f}) = \mathcal{F}_{s/f'}$. Note however that the relation $F_1(\mathcal{F}_{r/f}) = \mathcal{F}_{s/f'}$ is equivalent to $\rho(F_1(\mathcal{F}_{r/f})) = \rho(\mathcal{F}_{s/f'})$, or $\rho(F_1)(\rho(\mathcal{F}_{r/f})) = \rho(\mathcal{F}_{s/f'})$. Hence this number of edges is equal to the number of unordered pairs of faces $\mathcal{F}_{\gamma(r)/g}$ and $\mathcal{F}_{\gamma(s)/g'}$ for which $F_2(\mathcal{F}_{\gamma(r)/g}) = \mathcal{F}_{\gamma(s)/g'}$, i.e., the number of edges between vertices $\gamma(r)$ and $\gamma(s)$ in the face pairing graph for F_2 . Thus the face pairing graphs for F_1 and F_2 are isomorphic.

Conversely say face pairings F_1 and F_2 have isomorphic face pairing graphs G_1 and G_2 respectively. Let γ be an isomorphism from G_1 to G_2 , where γ maps vertices of G_1 to vertices of G_2 and edges of G_1 to edges of G_2 so that edge e meets vertex v if and only if edge $\gamma(e)$ meets vertex $\gamma(v)$.

Let S_1 and S_2 be the sets of faces $\mathcal{F}_{t/f}$ for which $F_1(\mathcal{F}_{t/f}) \neq \emptyset$ and $F_2(\mathcal{F}_{t/f}) \neq \emptyset$ respectively. From Definition 2.6.1 we see that γ maps unordered pairs of faces $\mathcal{F}_{t/f}, \mathcal{F}_{t'/f'}$ satisfying $F_1(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$ to unordered pairs of faces $\mathcal{F}_{\gamma(t)/g}, \mathcal{F}_{\gamma(t')/g'}$ satisfying $F_2(\mathcal{F}_{\gamma(t)/g}) = \mathcal{F}_{\gamma(t')/g'}$. We can thus define a map $\rho : S_1 \rightarrow S_2$ for which $\rho(\mathcal{F}_{t/f}) = \mathcal{F}_{\gamma(t)/g}$ and $\rho(\mathcal{F}_{t'/f'}) = \mathcal{F}_{\gamma(t')/g'}$ for each such unordered pair $\mathcal{F}_{t/f}, \mathcal{F}_{t'/f'}$.

Since γ is 1-to-1 and onto on both vertices and edges we see that $\rho : S_1 \rightarrow S_2$ is similarly 1-to-1 and onto, and that it maps every face of tetrahedron \mathcal{T}_t to a face of tetrahedron $\mathcal{T}_{\gamma(t)}$. We can thus extend ρ to a relabelling of order n by arbitrarily mapping faces $\mathcal{F}_{t/f} \in \mathfrak{F}_n \setminus S_1$ to faces $\mathcal{F}_{\gamma(t)/f} \in \mathfrak{F}_n \setminus S_2$.

Note that $F_1(\mathcal{F}_{t/f}) = \emptyset$ for each face $\mathcal{F}_{t/f} \in \mathfrak{F}_n \setminus S_1$ and that $F_2(\mathcal{F}_{t/f}) = \emptyset$ for each face $\mathcal{F}_{t/f} \in \mathfrak{F}_n \setminus S_2$. Letting $\rho(\emptyset) = \emptyset$ for convenience we can thus make the final observation that for all faces $\mathcal{F}_{t/f}$ we have $F_2(\rho(\mathcal{F}_{t/f})) = \rho(F_1(\mathcal{F}_{t/f}))$, regardless of whether $\mathcal{F}_{t/f} \in S_1$ or $\mathcal{F}_{t/f} \in \mathfrak{F}_n \setminus S_1$. Therefore $F_2 = \rho(F_1)$ and we see that face pairings F_1 and F_2 are indeed isomorphic. ■

Using Definition 2.6.1 we can translate the structural properties of face pairings into structural properties of their face pairing graphs, as seen in the following results.

Lemma 2.6.4 Every vertex of a face pairing graph has degree ≤ 4 . Furthermore, if a face pairing F has no faces mapping to \emptyset then every vertex of the corresponding face pairing graph has degree precisely 4.

Proof Consider vertex t of the face pairing graph for face pairing F . From Definition 2.6.1 we see that there is precisely one edge end meeting vertex t for each face $\mathcal{F}_{t/f}$ satisfying $\mathcal{F}_{t/f} \neq \emptyset$ (where each edge has two edge ends, and in particular an edge looping from vertex t back to itself provides two edge ends meeting vertex t).

The total number of edge ends meeting vertex t is the degree of vertex t , which is thus ≤ 4 since f can take precisely four possible values (0, 1, 2 and 3). Furthermore, it follows that the number of such edge ends is precisely 4 if and only if $F(\mathcal{F}_{t/f}) \neq \emptyset$ for all f . ■

Lemma 2.6.5 If G is a multigraph with n vertices all of which have degree ≤ 4 , then G is the face pairing graph for some face pairing F of order n . Furthermore, if every vertex of G has degree precisely 4 then F has no faces mapping to \emptyset .

Proof Arbitrarily number the vertices of G as $0, 1, \dots, n-1$. For each vertex i , arbitrarily label the edge ends meeting vertex i as $\mathcal{E}_{i/0}, \mathcal{E}_{i/1}, \dots, \mathcal{E}_{i/(d_i-1)}$ where $d_i \leq 4$ is the degree of vertex i . Then each edge will have two associated labels, one at each end. For each face $\mathcal{F}_{t/f}$ we now define $F(\mathcal{F}_{t/f}) = \mathcal{F}_{t'/f'}$ where labels $\mathcal{E}_{t/f}$ and $\mathcal{E}_{t'/f'}$ are at opposite ends of the same edge in G . If label $\mathcal{E}_{t/f}$ has not been assigned to an edge end, i.e., if $f \geq d_t$, then we define $F(\mathcal{F}_{t/f}) = \emptyset$.

It is straightforward to see that F satisfies the requirements of a face pairing of order n as described in Definition 2.1.2. Furthermore, if every vertex G has degree precisely 4 then we never have $f \geq d_t$ for any face $\mathcal{F}_{t/f}$ and so F has no faces mapping to \emptyset . ■

Lemma 2.6.6 A face pairing is connected if and only if its face pairing graph is connected.

Proof Recall from Definition 2.1.8 that a face pairing F is connected if, for any pair of tetrahedra \mathcal{T}_a and \mathcal{T}_b , there is some sequence of tetrahedra

$$\mathcal{T}_a = \mathcal{T}_{i_0}, \mathcal{T}_{i_1}, \mathcal{T}_{i_2}, \dots, \mathcal{T}_{i_k} = \mathcal{T}_b$$

in which each two consecutive tetrahedra $\mathcal{T}_{i_j}, \mathcal{T}_{i_{j+1}}$ are joined, i.e., there is some face $\mathcal{F}_{i_j/f}$ of tetrahedron \mathcal{T}_{i_j} for which $F(\mathcal{F}_{i_j/f})$ is a face of tetrahedron $\mathcal{T}_{i_{j+1}}$.

From Definition 2.6.1 we see that this is true precisely when for any graph vertices a and b there is some sequence of vertices

$$a = i_0, i_1, i_2, \dots, i_k = b$$

in which each two consecutive vertices i_j, i_{j+1} are joined by an edge. That is, face pairing F is connected if and only if its face pairing graph is connected. ■

It follows from the results above that the isomorphism classes of connected 4-valent graphs correspond precisely to the isomorphism classes of connected face pairings associated with triangulations that have no boundary faces. The set of all isomorphism classes of connected 4-valent graphs with 1, 2 and 3 vertices is illustrated in Figure 2.23.




Figure 2.23: All connected 4-valent graphs with 1, 2 and 3 vertices

2.6.2 Searching for Subgraphs

As promised in the opening remarks of Section 2.6, we begin now to develop tests that allow us to quickly discard face pairings that can never lead to closed orientable prime minimal triangulations.

All of these tests involve searching for specific undesirable subgraphs within face pairing graphs.

Theorem 2.6.7 Let F be a face pairing of order $n \geq 3$. If the face pairing graph for F contains a *triple edge* (two vertices joined by three distinct edges, as illustrated in Figure 2.24), then F cannot be the associated face pairing of a closed orientable prime minimal triangulation.




Figure 2.24: A triple edge in a face pairing graph

Proof Suppose that T is a closed orientable prime minimal triangulation with ≥ 3 tetrahedra whose face pairing graph contains a triple edge. Observe that this triple edge corresponds to two distinct tetrahedra of T that are joined along three different faces. This proof enumerates all possible ways in which this can be done and derives a contradiction in each case.

Let these two tetrahedra be $ABCD$ and $A'B'C'D'$ as illustrated in Figure 2.25. Faces ABD and $A'B'D'$ are identified, faces BCD and $B'C'D'$ are identified and faces CAD and $C'A'D'$ are identified, though not necessarily with these specific vertex identifications; faces may be rotated or reflected before they are identified.




Figure 2.25: Two tetrahedra to be joined along three faces

Faces ABC and $A'B'C'$ remain unaccounted for, though they cannot be identified with each other since this would produce a 2-tetrahedron triangulation. Thus faces ABC and $A'B'C'$ are distinct faces of T ; we refer to these as the *boundary faces* since they bound the subcomplex formed by the two tetrahedra under investigation.

Each specific method of joining our two tetrahedra along the faces described is denoted by a *matching string*. A matching string is a sequence of three symbols representing the transformations that are applied to faces ABD , BCD and CAD respectively before they are identified with their counterparts from the other tetrahedron. Each symbol is one of the following.

- ι : No transformation is applied.
- κ : The face is rotated clockwise.
- α : The face is rotated anticlockwise.
- c : The face is reflected so that the centre point of the diagram (i.e., point D) remains fixed.
- l : The face is reflected so that the point at the clockwise end of the face (e.g., point B on face ABD) remains fixed.

- r : The face is reflected so that the point at the anticlockwise end of the face (e.g., point A on face ABD) remains fixed.

The full list of exact face identifications corresponding to each symbol is given in Table 2.3.

	ι	κ	α	c	l	r
ABD	$A'B'D'$	$B'D'A'$	$D'A'B'$	$B'A'D'$	$D'B'A'$	$A'D'B'$
BCD	$B'C'D'$	$C'D'B'$	$D'B'C'$	$C'B'D'$	$D'C'B'$	$B'D'C'$
CAD	$C'A'D'$	$A'D'C'$	$D'C'A'$	$A'C'D'$	$D'A'C'$	$C'D'A'$

Table 2.3: Precise face identifications corresponding to each transformation symbol

We can then list the matching strings for all possible ways of joining our two tetrahedra along these three face pairs. Each matching string is listed only once up to equivalence, where equivalence includes rotating the two tetrahedra, reflecting the two tetrahedra and swapping the two tetrahedra. Note also that since T is orientable, a matching string cannot contain both a symbol from $\{\iota, \kappa, \alpha\}$ and a symbol from $\{c, l, r\}$. The final list of matching strings is as follows.

$$\iota\iota, \iota\kappa, \iota\kappa\kappa, \iota\kappa\alpha, \kappa\kappa\kappa, \kappa\kappa\alpha, ccc, ccl, cl, clr, clr, lll, llr.$$

Note that matching strings $\iota\kappa\alpha$ and $\iota\alpha\kappa$ are equivalent by swapping the two tetrahedra, but clr and crl are not.

These thirteen types of matching can be split into three categories, where the matchings in each category give rise to a similar contradiction using almost identical arguments. We examine each category in turn.

- **Degree two edges ($\iota\iota, \iota\kappa, cl, llr$):** Simply by following the edge identifications induced by our chosen face identifications, we can see that each of these matchings gives rise to a non-boundary edge of degree two, contradicting Theorem 2.3.6.

An example of this is illustrated in Figure 2.26 which shows the induced edge identifications for the matching llr . The non-boundary edge of degree two here is edge DA (or equivalently $A'D'$).




Figure 2.26: Edge identifications for the matching llr

- **Spherical subcomplexes ($\iota\kappa\kappa, \iota\kappa\alpha, \kappa\kappa\kappa, \kappa\kappa\alpha, clr$):** Again we follow the induced edge identifications, but in this case the consequence is that we see the two boundary faces are joined at their edges to form a two-triangle sphere. This contradicts either Theorem 2.4.15 or Theorem 2.4.16 according to which specific triangulation of the 2-sphere has been produced. This behaviour is illustrated in Figure 2.27 for the matching $\kappa\kappa\alpha$. Here faces ABC and $A'B'C'$ form a sphere as described in Theorem 2.4.16.




Figure 2.27: Edge identifications for the matching $\kappa\kappa\kappa$

Similarly, matchings $\iota\kappa\kappa$, $\iota\kappa\alpha$ and clr produce a sphere as described in Theorem 2.4.16, whereas matching $\kappa\kappa\kappa$ produces a sphere as described in Theorem 2.4.15.

- **Bad vertex links (ccc , ccl , cll , lll):** Consider matching ccc . The induced edge and vertex identifications are illustrated in the left hand portion of Figure 2.28; in particular we see that vertices D and D' are identified with each other but with none of A , B , C , A' , B' or C' . The link of vertex D (or equivalently D') is formed from the two triangles shown on the right hand side of Figure 2.28, which join together to form a torus. Thus in the overall triangulation this vertex must form a torus cusp and we cannot have a triangulation of a closed 3-manifold.




Figure 2.28: The vertex link of D and D' for the matching ccc

For matchings ccl , cll and lll , all eight vertices A , B , C , D , A' , B' , C' and D' are identified as a single vertex in the triangulation. In particular, this vertex lies on the boundary faces ABC and $A'B'C'$, and so the vertex link as restricted to our two tetrahedra will be incomplete (i.e., will have boundary components). However, for all three of these matchings, this (partial) vertex link is observed to be a once-punctured torus. Thus however the entire triangulation T is formed, the complete link of our vertex in T cannot be a sphere (since there is no way to fill in the boundary of a punctured torus to form a sphere). So again we have a contradiction since T cannot be a triangulation of a closed 3-manifold.

As an example of this behaviour, Figure 2.29 illustrates the induced edge and vertex identifications for matching lll as well as the corresponding vertex link (shown as eight individual triangles followed by a combined figure) which we see is indeed a punctured torus.

Thus we see that every method of identifying three faces of the first tetrahedron with three faces of the second gives rise to a contradiction, and so our result is established. ■

Definitions 2.6.8 A *chain of length k* is the multigraph formed as follows. Take $k + 1$ vertices labelled $0, 1, 2, \dots, k$ and join vertices i and $i + 1$ with a double edge for all $0 \leq i \leq k - 1$. Each of these edges is called an *interior edge* of the chain.




Figure 2.29: The partial vertex link for the matching lll

If a loop is added joining vertex 0 to itself the chain becomes a *one-ended chain*. If another loop is added joining vertex k to itself the chain becomes a *double-ended chain* (and is now a 4-valent multigraph). These loops are called *end edges* of the chain.

Example 2.6.9 A one-ended chain of length 4 is illustrated in Figure 2.30, and a double-ended chain of length 3 is illustrated in Figure 2.31.



Figure 2.30: A one-ended chain in a face pairing graph



Figure 2.31: A double-ended chain in a face pairing graph

Lemma 2.6.10 Let T be a closed orientable prime minimal triangulation with ≥ 3 tetrahedra and associated face pairing graph G . If G contains a one-ended chain then the tetrahedra of T corresponding to the vertices of this one-ended chain form a layered solid torus in T , as described by Definition 1.2.2.

Proof We prove this by induction on the chain length. A one-ended chain of length 0 consists of a single vertex with a single end edge, representing a single tetrahedron with two of its faces identified. If these faces are simply snapped shut as illustrated in the left hand diagram of Figure 2.32 (faces ABC and DBC being identified), the edge between them will have degree one in the final triangulation which cannot happen according to Theorem 2.3.7. Thus these faces are identified with a twist as illustrated in the right hand diagram of Figure 2.32 (faces ABC and BCD being identified), producing the layered solid torus $LST(1, 2, -3)$.

Assume then that any one-ended chain of length k must correspond to a layered solid torus (which will have $k + 1$ tetrahedra), and consider a one-ended chain of length $k + 1$. This one-ended chain is simply a one-ended chain of length k with an extra double edge attached to the end, and so by our inductive hypothesis the corresponding $k + 2$ tetrahedra must form a $(k + 1)$ -tetrahedron layered solid torus with an additional tetrahedron joined to its boundary along two faces.




Figure 2.32: Identifying two faces of a tetrahedron

By symmetry of the two-triangle torus (which forms the layered solid torus boundary), we can picture the situation as shown in the left hand diagram of Figure 2.33, where face WXY of the new tetrahedron is to be joined directly to face ABC of the layered solid torus, and face YZW of the new tetrahedron is to be joined to face CDA but possibly with the vertices identified in some different order (i.e., with the faces being rotated before they are identified).




Figure 2.33: Appending a new tetrahedron to a layered solid torus

If face YZW is simply folded over onto face CDA , i.e., the vertices are identified in this precise order with no rotation, then we have simply layered the new tetrahedron onto edge AC , and so we obtain a larger layered solid torus (with $k + 2$ tetrahedra) as required.

If on the other hand a rotation takes place, we may by symmetry assume that face YZW is identified with face ACD , as illustrated in the right hand diagram of Figure 2.33. In this diagram the two new boundary faces are shaded (XWZ and XYZ), and we see that their edges are in fact identified to form a 2-sphere (which may be pinched at vertices W , X , Y and Z). These new boundary faces thus satisfy the conditions of Theorem 2.4.16, and we see that T cannot be both prime and minimal. ■

Theorem 2.6.11 Let F be a face pairing of order $n \geq 3$. If the face pairing graph for F contains as a subgraph a *broken double-ended chain* (i.e., a double-ended chain missing one interior edge, as illustrated in Figure 2.34) and this face pairing graph is not simply a double-ended chain itself, then F cannot be the associated face pairing of a closed orientable prime minimal triangulation.




Figure 2.34: A broken double-ended chain in a face pairing graph

Proof Observe that a broken double-ended chain is merely a pair of one-ended chains joined by an edge. Assume T is a closed orientable prime minimal triangulation whose face pairing graph contains a broken double-ended chain; then Lemma 2.6.10 implies that T contains a pair of layered solid tori whose boundaries are joined along one face.

This situation is depicted in the left hand diagram of Figure 2.35, where the two torus boundaries of the layered solid tori are shown and where face ABC is identified with face XYZ . The resulting edge identifications of the remaining two boundary faces are illustrated in the right hand diagram of this figure; in particular, it can be seen that these remaining boundary faces form a two-triangle sphere (though with all three vertices pinched together, since each layered solid torus has only one vertex).




Figure 2.35: Two layered solid tori joined along a face

Let these remaining boundary faces be F_1 and F_2 . If these faces are not identified, then F_1 and F_2 satisfy the conditions of Theorem 2.4.15 and so T is either non-prime or non-minimal. Since we know this is not true, faces F_1 and F_2 must be identified. But then in the face pairing graph for T , the single edge between the two one-ended chains now becomes a double edge and we see that the graph contains an entire double-ended chain.

Finally, since every vertex in a double-ended chain has degree 4, Lemma 2.6.4 implies that the face pairing graph is *only* this double-ended chain, contradicting the conditions of our theorem. ■

The proof of the next result regarding unallowable subgraphs (Theorem 2.6.13) is a massive case analysis. In the interests of making this case analysis more bearable we present the following result relating to double edges in a face pairing graph.

Lemma 2.6.12 Let T be a closed orientable prime minimal triangulation with ≥ 3 tetrahedra. If two distinct tetrahedra of T are joined to each other along two distinct faces then these two face identifications must be as illustrated in one of the diagrams of Figure 2.36.




Figure 2.36: Allowable ways of joining two tetrahedra along two faces

Specifically, let the tetrahedra be $ABCD$ and $A'B'C'D'$, with faces ABC and $A'B'C'$ identified and with faces BCD and $B'C'D'$ identified (though not necessarily with these precise vertex identifications; the faces may be rotated or reflected before they are identified).

Then, allowing for the two tetrahedra to be relabelled and/or swapped, this result states that we must have one of the following two cases. In the first case, face ABC is identified with $A'B'C'$ (i.e., no rotations or reflections take place) and face BCD is identified with $C'D'B'$ (i.e., the faces are rotated before being identified); this is illustrated in the left hand diagram of Figure 2.36. In the second case, face ABC is identified with $C'A'B'$ and face BCD is identified with $C'D'B'$ (i.e., both identifications involve a rotation in the same direction); this is illustrated in the right hand diagram of Figure 2.36.

Proof Again allowing for the tetrahedra to be relabelled and/or swapped, the only ways of identifying two tetrahedra of T along two distinct faces that are not listed as acceptable according to this lemma are the two methods illustrated in Figure 2.37.

In the first method, face ABC is identified with $A'B'C'$ and face BCD is identified with $B'C'D'$ (i.e., no rotations or reflections take place for either identification). These identifications are illustrated in the left hand diagram of Figure 2.37. In the second method, face ABC is identified with $B'C'A'$ and face BCD is identified with $C'D'B'$ (i.e., both identifications involve a rotation but in opposite directions). These identifications are illustrated in the right hand diagram of Figure 2.37.




Figure 2.37: Unallowable ways of joining two tetrahedra along two faces

The first method is easily disposed of since edge BC (identified with $B'C'$) becomes a non-boundary edge of degree two, contradicting Theorem 2.3.6. The second method requires a little more work.




Figure 2.38: A pinched ball formed from two tetrahedra

The left hand diagram of Figure 2.38 illustrates the structure formed by identifying these two tetrahedra using the second method. The structure is a ball, with two opposite points (vertices B and C') pinched together at the centre. Let edge e denote the edge forming the upper rim of the structure, as illustrated in the right hand diagram of Figure 2.38 (this is in fact edge AD).

Then edge e bounds a disc whose interior is contained entirely within the interior of the pinched ball, with the single exception of the pinched vertex at the ball's centre. The dashed curves in the right hand diagram of Figure 2.38 illustrate where this disc lies. Since the interior of this disc lies within the interior of the ball, the only possible way in which the disc might intersect itself is if the two vertices marked in the right hand diagram of Figure 2.38 are identified, i.e., the vertex at the centre of the pinched ball is identified with the vertex forming the endpoints of edge e .

However, if these two vertices are identified then this self-intersection takes place only at a single point, and so we can remove this self-intersection by shifting the disc interior slightly away from the disc boundary in the underlying 3-manifold, as illustrated in Figure 2.39 (note that this will move a small portion of the disc interior outside the pinched ball). Thus we can manipulate the interior of this disc to show that edge e in fact bounds an embedded disc in the underlying 3-manifold.




Figure 2.39: Removing the self-intersection of the disc bounded by e

We may now call upon Jaco and Rubinstein's 0-efficiency results. Lemma 2.4.10 implies that our triangulation is not 0-efficient or is a triangulation of S^3 . Meanwhile, Corollary 2.4.7 shows that since our triangulation is minimal it must be 0-efficient, and since S^3 can be formed from a single tetrahedron but we have ≥ 3 tetrahedra then our triangulation cannot be of S^3 . So we again have a contradiction. ■

Theorem 2.6.13 Let F be a face pairing of order $n \geq 3$. If the face pairing graph for F contains as a subgraph a *one-ended chain with a double handle* (a double-ended chain with one end edge replaced by a triangle containing one double edge, as illustrated in Figure 2.40), then F cannot be the associated face pairing of a closed orientable prime minimal triangulation.




Figure 2.40: A one-ended chain with a double handle in a face pairing graph

Proof Let T be a closed orientable prime minimal triangulation with ≥ 3 tetrahedra whose face pairing graph contains a one-ended chain with a double handle. From Lemma 2.6.10 we see that the one-ended chain corresponds to a layered solid torus in T . The double handle in turn must correspond to two additional tetrahedra each of which is joined to the other along two faces and each of which is joined to one of the boundary faces of the layered solid torus.

This construction is illustrated in Figure 2.41. The layered solid torus is below faces ABC and BCD which form the boundary of this layered solid torus. The two additional tetrahedra are $XABC$ and $YBCD$; observe that each of these tetrahedra is joined to one of the boundary faces




Figure 2.41: Tetrahedra corresponding to a one-ended chain with a double handle

of the layered solid torus. The two new tetrahedra are then joined to each other along two faces; specifically faces XAC and YCB are identified and faces XCB and DCY are identified.

As with the proof of Theorem 2.6.7, our proof will involve enumerating all possible ways in which this construction can be carried out and in each case deriving a contradiction. To assist in this task we will describe a simple way of representing each possible variant of the above construction.

The scenario described in Figure 2.41 can be distilled into a simplified diagram as illustrated in Figure 2.42. We begin with the two-triangle torus that forms the boundary of the layered solid torus as shown in the left hand diagram of Figure 2.42. We then add our two new tetrahedra, converting the two-triangle torus into a six-triangle torus as illustrated in the central diagram of Figure 2.42. Finally we add markings to the diagram to illustrate how the two new tetrahedra are to be joined along two faces. This is done by marking vertices a, b and c of the first face and vertices x, y and z of the second face along which the two tetrahedra are joined. It can be seen from the right hand diagram of Figure 2.42 that faces XAC and YCB are identified and faces XCB and DCY are identified as described earlier. For clarity the two faces that are not yet identified with any other face are shaded.




Figure 2.42: A simplified diagram for a one-ended chain with a double handle

In a similar fashion we can represent any set of tetrahedra corresponding to a one-ended chain with a double handle using a diagram similar to the right hand diagram of Figure 2.42. In this way we will enumerate all possible diagrams and in each case prove that T cannot be a closed orientable prime minimal triangulation as claimed.

Recall from Lemma 2.6.12 that if T is a closed orientable prime minimal triangulation then there are restrictions upon the possible ways in which our two new tetrahedra can be joined along two faces. By ignoring all diagrams that do not conform to Lemma 2.6.12 and by exploiting the symmetry of the two-triangle torus that forms the layered solid torus boundary we can reduce the set of all possible diagrams to those depicted in Figure 2.43.

As in the proof of Theorem 2.6.7 we can divide our eleven different diagrams $\alpha_1, \dots, \delta_2$ into a




Figure 2.43: All diagrams for a one-ended chain with a double handle

small number of categories, where the diagrams in each category give rise to a similar contradiction using almost identical arguments. To assist with this process the edge identifications induced in each diagram by the corresponding face identifications are shown in Figure 2.44. The diagrams can then be split into categories as follows.




Figure 2.44: Edge identifications corresponding to the different face identification diagrams

- **Spherical subcomplexes ($\alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3, \delta_1$):** In each of these diagrams we can observe that the two shaded faces are joined at their edges to form a two-triangle sphere. Specifically diagrams $\alpha_1, \alpha_2, \beta_1$ and β_3 produce spheres as described in Theorem 2.4.16 and diagrams α_3, β_2 and δ_1 produce spheres as described in Theorem 2.4.15. In each case Theorem 2.4.16 or Theorem 2.4.15 contradicts the claim that T is a closed orientable prime minimal triangulation.

- **Bad vertex links** ($\gamma_1, \gamma_2, \gamma_3, \delta_2$): In each of these diagrams it can be observed that all six vertices illustrated are in fact identified as a single vertex in the triangulation. We can calculate the link of this vertex as restricted to the portion of the triangulation that we are examining; for diagrams γ_1, γ_2 and γ_3 this link is observed to be a once-punctured torus and for diagram δ_2 it is observed to be a once-punctured two-holed torus. In each case however the remainder of triangulation T is constructed it is impossible for this vertex link to be extended to become a sphere. Thus T cannot be a triangulation of a closed 3-manifold.

The vertex link calculation is illustrated in Figure 2.45 for diagram γ_1 . The disc on the left with edges p, q, r and s represents the vertex link of the layered solid torus and the triangles beside it represent the pieces of vertex link taken from the two new tetrahedra. These pieces are combined into a single surface on the right hand side of the diagram which we see is indeed a once-punctured torus.




Figure 2.45: Calculating the vertex link for diagram γ_1

Thus for each diagram of Figure 2.43 we observe that T cannot be a closed orientable prime minimal triangulation, concluding our proof. ■

2.6.3 Face Pairing Statistics

How useful then are the results of Section 2.6.2? As a first estimate, Table 2.4 lists how many face pairing graphs on n vertices for $3 \leq n \leq 11$ contain each of the undesirable subgraphs described in Section 2.6.2 (a triple edge, a broken double-ended chain or a one-ended chain with a double handle). Only connected 4-valent graphs, i.e., graphs whose corresponding triangulations are connected with no boundary faces, are considered. The results presented were obtained using the program *Regina*.

All counts in this table are given up to isomorphism, so that if graph G has already been counted then no other graph isomorphic to G will be counted. This style of counting allows us to directly relate these totals to the census algorithm in which we only generate one face pairing from each isomorphism class. The meanings of the individual columns of Table 2.4 are as follows.

- *Vertices*: The number of vertices n .
- *Total*: The total number t of connected 4-valent graphs with n vertices.
- *None*: The number and percentage of graphs from this total t in which none of the undesirable subgraphs listed above were found.

- *Some*: The number of graphs from this total t in which at least one of the undesirable subgraphs listed above was found.
- *Triple*: The number of graphs from the total t that contain a triple edge.
- *Broken*: The number of graphs from the total t that contain a broken double-ended chain.
- *Handle*: The number of graphs from the total t that contain a one-ended chain with a double handle.
- *Time*: The running time taken to calculate the values in this row of the table. Running times are measured on a single 1.2GHz Pentium III processor and are displayed as *h:mm:ss*.

Note that the *Total* column should equal the *None* column plus the *Some* column for each number of vertices. Note also that the sum of the *Triple*, *Broken* and *Handle* columns might exceed the *Some* column since some graphs may contain more than one type of undesirable subgraph.

Vertices	Total	None	Some	Triple	Broken	Handle	Time
3	4	2 (50%)	2	1	1	1	0:00
4	10	4 (40%)	6	3	3	2	0:00
5	28	12 (43%)	16	8	10	4	0:00
6	97	39 (40%)	58	29	36	12	0:00
7	359	138 (38%)	221	109	137	40	0:01
8	1 635	638 (39%)	997	497	608	155	0:05
9	8 296	3 366 (41%)	4 930	2 479	2 976	685	0:44
10	48 432	20 751 (43%)	27 681	14 101	16 568	3 396	7:21
11	316 520	143 829 (45%)	172 691	88 662	102 498	18 974	1:20:48

Table 2.4: Frequency of undesirable structures within face pairing graphs

As a reminder of just how fast face pairing generation is within the overall census algorithm, observe from Table 2.4 that we can construct all 316 520 face pairings of order 11 and eliminate over half of them in under an hour and a half, compared to over 27 months of running time required to complete an entire orientable census for only 8 tetrahedra as seen in Table 2.1 on page 65.

Examining Table 2.4 we see then that for each number of vertices a little over half of the possible face pairing graphs can be eliminated using the results of Section 2.6.2, assuming we are interested only in closed orientable prime minimal triangulations. Whilst this does not reduce the complexity of our algorithm's running time, it does allow us to reduce the running time by more than 50%, a significant improvement for a census that may take months or years to complete.

Beyond the number of face pairings graphs that are eliminated however, we are also interested in the number of face pairing graphs that cannot lead to a closed orientable prime minimal triangulation but that are not identified as such by the results of Section 2.6.2. This allows us to estimate how much potential there is for improvement of these results and hence further improvement of the running time of the census algorithm. We will examine in detail the sets of all face pairing graphs with 3, 4 and 5 vertices.

For the 3-vertex face pairing graphs listed in Figure 2.23, one of these graphs has a triple edge, one has a broken double-ended chain and the other two both lead to closed orientable prime minimal triangulations with three tetrahedra.

A list of all connected 4-valent 4-vertex graphs is shown in Figure 2.46. The graphs to the left of the dotted line all lead to closed orientable prime minimal triangulations; the graphs to the right of the dotted line do not. It can be seen that all of the right hand graphs contain either a triple edge, a broken double-ended chain or a one-ended chain with a double handle.




Figure 2.46: All connected 4-valent graphs with 4 vertices

Thus for three and four tetrahedra, the results of Section 2.6.2 in fact perfectly divide the face pairings into those that lead to desirable triangulations and those that do not.

For five tetrahedra these theorems no longer perfectly divide the face pairings as we would like. Figure 2.47 lists all 5-vertex face pairing graphs; there are 28 in total. The 8 graphs in the top section all lead to closed orientable prime minimal triangulations and the remaining 20 do not. Of these remaining 20, the 16 graphs in the middle section each contain a triple edge, a broken double-ended chain or a one-ended chain with a double handle. We see then that the final 4 graphs in the bottom section can never lead to a desirable triangulation but are not identified as such by the theorems in Section 2.6.2.




Figure 2.47: All connected 4-valent graphs with 5 vertices

2.7 Future Directions

Although the census results described in Chapter 3 extend no further than six tetrahedra, the initial generation of closed orientable triangulations of seven and eight tetrahedra has already been

carried out as seen in Section 2.5. It is desirable to extend these census results to nine tetrahedra and beyond, but since the eight tetrahedron census consumed 27 months of processor time it is clear that improvements to the census algorithm itself will need to be made.

2.7.1 Eliminating Face Pairings

The observations of Section 2.6.3 certainly offer a promising start to improving the census algorithm. Using Theorems 2.6.7, 2.6.11 and 2.6.13 we can refuse to process any face pairing whose face pairing graph contains a triple edge, a broken double-ended chain or a one-ended chain with a double handle. Table 2.4 suggests that this should reduce the overall running time of our algorithm by over 50%. These algorithmic modifications have already been implemented in *Regina*.

Furthermore, from our analysis of all 4-valent face pairing graphs with 5 vertices we see that there are face pairings that do not lead to closed orientable prime minimal triangulations but that are not covered by these three theorems. One direction for future research is therefore to prove additional results of a similar nature that can be used to identify further face pairings that do not need to be processed, thus allowing for a further reduction of the running time of the algorithm.

2.7.2 Improved Gluing Permutation Generation

Consider once more Algorithm 2.1.16. This describes the census algorithm as a three-stage process in which we generate face pairings, generate gluing permutation selections for each face pairing and then construct and analyse the corresponding triangulations. Recall from Section 2.6.3 that the running time spent generating face pairings is inconsequential in the context of the entire census algorithm. In fact almost the entire running time of the algorithm is spent in the second stage generating gluing permutations.

As seen in Algorithm 2.3.2 this gluing permutation generation is a recursive process in which we choose from six permutations (or three for an orientable triangulation) for each pair of faces that are to be joined. This allows for up to 6^{2n} (or 3^{2n} for an orientable triangulation) different gluing permutation selections for any given face pairing.

It is thus in the generation of gluing permutation selections that we should seek to make the strongest improvements to the census algorithm. Section 2.3.1 already offers techniques with which this recursion can be pruned and thus made faster. It is possible that with further research we can devise additional constraints that allow us to undertake heavier pruning of this recursion and thus make the census algorithm more efficient.

It is possible however that instead of chipping away at the recursion with more and more intricate pruning techniques, we could perhaps make more significant improvements by substantially redesigning the recursion using results from Section 2.6.2. Note that these results can only be used in a census in which only closed orientable prime minimal triangulations are required.

Recall from Lemma 2.6.10 that every one-ended chain in a face pairing graph must correspond to a layered solid torus in the resulting triangulation. Recall also from Lemma 2.6.12 that every double edge in a face pairing graph must correspond to one of a restricted set of identifications of two tetrahedra along two faces. Instead of simply selecting permutations for each pair of faces one after another and pruning where possible as in Algorithm 2.3.2, we can thus redesign the entire recursion as follows.

1. Identify all one-ended chains in the underlying face pairing graph. Recursively select a layered solid torus of the appropriate size to correspond to each of these one-ended chains.
2. Identify all double edges in the underlying face pairing graph that do not belong to the one-ended chains previously processed. For each such double edge recursively select an identification of the two corresponding tetrahedra that conforms to Lemma 2.6.12.
3. At this point many of the individual permutations in the overall gluing permutation selection have already been established. Recursively select the remaining permutations as in the original algorithm, pruning where possible.

The first step of this redesigned algorithm should offer a substantial improvement in running time, as can be seen by the following rough calculations. In a one-ended chain on k vertices there are $2k - 1$ edges corresponding to $2k - 1$ permutations that must be selected. Using the original algorithm we may be investigating up to $3^{2k-1} = \frac{1}{3}9^k$ possible sets of permutations for these edges, although this number will be reduced due to pruning.

Using the redesigned algorithm however, instead of counting all possible sets of permutations we need only count the number of possible layered solid tori for this one-ended chain. In constructing such a layered solid torus we have two choices for how the base tetrahedron of the layered solid torus is joined to itself and then two choices for how each subsequent tetrahedron is layered on (there are in fact three choices for each layering but one will always lead to an edge of degree two which from Corollary 2.3.8 can be ignored). Thus there are only 2^k possible layered solid tori that can correspond to this one-ended chain, a vast improvement upon the original estimate of $\frac{1}{3}9^k$ possible sets of permutations.

The redesigned algorithm as described above is yet to be implemented in *Regina*, although this will be done in the near future. Work is also continuing to obtain further results in a similar vein to Lemmas 2.6.10 and 2.6.12, in the hope that face pairing graphs can be used to impose an even stronger structure upon the resulting triangulations. Such results if obtained could be similarly incorporated into the generation of gluing permutations in the hope of further streamlining the census algorithm.

Chapter 3

Minimal Triangulations

In this chapter we harvest the results of the census whose underlying algorithm is described in Chapter 2. In particular we present all closed prime minimal triangulations formed from at most six tetrahedra, including triangulations of both orientable and non-orientable 3-manifolds. The main body of this chapter is devoted to an extensive analysis of the combinatorial structures found within these triangulations.

The highlight of this chapter is the categorisation of most of these triangulations into a small number of parameterised families. Each of these families represents an infinite class of triangulations sharing a common large-scale structure. Of the 218 different triangulations presented, 204 of these triangulations can be represented as members of these families. Of the remaining 14 triangulations, four are exceptional triangulations of simple 3-manifolds that use very few tetrahedra and are not expected to generalise. The final ten triangulations have yet to be categorised although their underlying 3-manifolds are understood.

Such a categorisation of triangulations into infinite families is certainly appealing. Large classes of triangulations may be simultaneously studied, and algorithms are easily developed for generating triangulations of infinite classes of 3-manifolds.

Furthermore, when presented with an arbitrary triangulation of an unknown 3-manifold, having a rich collection of such families at our disposal increases the chance that we can manipulate the triangulation into a form that is recognisable as a member of one of these families. Recognising a triangulation in this way then allows the underlying 3-manifold to be established. As discussed in Section 2.4.4, this ability to recognise previously unseen triangulations of unknown 3-manifolds is a powerful tool in the processing of a newly formed census of triangulations. An example of its application is seen in Section 4.5 in which we identify the underlying 3-manifolds of seven-tetrahedron and eight-tetrahedron triangulations containing particular types of splitting surfaces.

The results of this chapter are an extension of the work of Matveev, Martelli and Petronio. In [29] Matveev presents all closed orientable prime minimal triangulations containing at most six tetrahedra, and in addition he describes a small number of families of triangulations. His constructions although presented in terms of special spines of 3-manifolds are analogous to layered lens spaces, twisted layered loops and augmented triangular solid tori as discussed in Section 3.3.

Martelli and Petronio then present the results of a nine-tetrahedron orientable census in [24]. Whilst this census extends to a larger number of tetrahedra, the focus of [24] is upon the underlying

3-manifolds and not the combinatorial structures of their minimal triangulations. These authors nevertheless extend the families of Matveev to include constructions analogous to layered chain pairs and a restricted class of plugged triangular solid tori as again discussed in Section 3.3.

It has come to the author's attention that during the final phases of preparing the material presented here, Martelli and Petronio have independently released further results in [26] that describe in more detail the combinatorial structures of some minimal triangulations. Their interest however remains in the underlying 3-manifolds, and as such their results focus upon finding any minimal triangulation of a given 3-manifold. In contrast to this approach, the results presented in this chapter analyse all minimal triangulations of each underlying 3-manifold.

As a result of the focus upon the combinatorial structures of all minimal triangulations, this chapter examines additional classes of triangulations that are not discussed by the above authors. The families of triangulations in this chapter are furthermore presented as direct constructions of triangulations built from tetrahedra, in contrast to the constructions presented by Matveev, Martelli and Petronio which are phrased in the context of special spines of 3-manifolds.

Of particular interest is the non-orientable census described in Section 3.5, whose results are not touched upon by any of the aforementioned authors. Of interest also is the enumeration of all vertex normal surfaces of each triangulation as presented in Section 3.6.

In Section 3.1 we outline the notation that is used for 3-manifolds and other structures throughout this chapter. Section 3.2 then presents a series of medium-sized building blocks from which our families of triangulations are constructed. Each of these families of triangulations is described in detail in Section 3.3. Sections 3.4 and 3.5 present the final results of the orientable and non-orientable censuses respectively, and Section 3.6 describes the enumeration of all vertex normal surfaces of these triangulations. Possible directions for further research are then outlined in Section 3.7.

Both the generation of triangulations within each census and the resulting analysis of these triangulations were performed using the program *Regina* as described in Section 1.3.

3.1 Notation

We begin by outlining the notation that is used throughout this chapter to represent the different topological spaces that are encountered.

3.1.1 1-Manifolds and 2-Manifolds

The symbols S^1 and I as usual represent the circle and the unit interval $[0, 1]$ respectively. In addition the following symbols represent the following frequently used surfaces.

T^2 : the torus	B^2 : the disc (i.e., the 2-ball)
$\mathbb{R}P^2$: the projective plane	M^2 : the Möbius band
K^2 : the Klein bottle	

3.1.2 Seifert Fibred Spaces

The notation used throughout this chapter for orientable Seifert fibred spaces is the same as is used by Matveev [29] as well as Martelli and Petronio in their earlier works [24]. An orientable

Seifert fibred space is represented as

$$\text{SFS}(F : (\alpha_1, \beta_1), (\alpha_2, \beta_2), \dots, (\alpha_k, \beta_k)),$$

where F describes the orbit manifold of the Seifert fibred space and the (α_i, β_i) pairs describe both the exceptional fibres and the additional twisting parameter used by other authors such as Orlik [32]. For each i the integers α_i and β_i are coprime and α_i is strictly positive.

The space $\text{SFS}(F : (\alpha_1, \beta_1), (\alpha_2, \beta_2), \dots, (\alpha_k, \beta_k))$ is formed by constructing the product space $F \times S^1$ (or the orientable twisted product $F \tilde{\times} S^1$ if F is non-orientable) and performing a drilling operation for each (α_i, β_i) pair. Each such operation involves drilling out a trivially fibred solid torus and then inserting a new solid torus whose meridional disc is bounded by the curve $\alpha_i\mu + \beta_i\lambda$, where μ represents a meridinal curve of the original torus that was drilled out and λ represents one of the circular fibres of the product $F \times S^1$. If $\alpha_i > 1$ then this operation will produce an exceptional fibre of index α_i .

Such a representation of a Seifert fibred space is not unique. In particular, the following results can be shown.

- The insertion of an additional $(1, 0)$ pair into the representation of a Seifert fibred space has no effect upon the underlying 3-manifold or its fibration;
- Two pairs (α_i, β_i) and (α_j, β_j) can be replaced with the pairs $(\alpha_i, \beta_i + t\alpha_i)$ and $(\alpha_j, \beta_j - t\alpha_j)$ without changing the underlying 3-manifold or its fibration for any integer t ;
- A collection of pairs $(1, t_1), (1, t_2), \dots, (1, t_m)$ can be replaced with the single pair $(1, t_1 + t_2 + \dots + t_m)$ without changing the underlying 3-manifold or its fibration;
- Every pair (α_i, β_i) can be simultaneously replaced with the pair $(\alpha_i, -\beta_i)$ without changing the underlying 3-manifold or its fibration.

The notation described above was chosen because the absence of bounds upon the parameters β_i allows for a more convenient parameterised description of the Seifert fibred spaces encountered amongst the families of triangulations discussed in Section 3.3.

Other authors including Orlik [32] and Martelli and Petronio in their later works [26] use a normalised form of this notation in which each (α_i, β_i) pair satisfies $0 < \beta_i < \alpha_i$ and describes precisely one exceptional fibre. In addition they include a single twisting parameter t which corresponds to a single $(1, t)$ pair in our notation. The equivalences listed above allow for a straightforward conversion between the two styles of representation.

3.1.3 Orbit Spaces

Although every closed orientable triangulation presented in this chapter represents a Seifert fibred space, notation of the form $\text{SFS}(F : (\alpha_1, \beta_1), (\alpha_2, \beta_2), \dots, (\alpha_k, \beta_k))$ is a little unwieldy. In many cases we can refer to a 3-manifold in a more compact fashion. We therefore introduce the following notation which is the same notation used by Matveev in [29].

The space S^3/G represents the orientable orbit space of the finite group G acting freely and orthogonally on the 3-sphere. It is shown by Hopf [13] and Milnor [31] that any such group G must be of one of the following forms.

- The cyclic group \mathbb{Z}_n ;
- The group Q_{4n} with presentation

$$Q_{4n} = \langle x, y : x^2 = (xy)^2 = y^n \rangle;$$

- The group $D_{2^k(2n+1)}$ with presentation

$$D_{2^k(2n+1)} = \langle x, y : x^{2^k} = 1, y^{2n+1} = 1, xy = y^{-1}x \rangle;$$

- The binary tetrahedral group P_{24} , the binary octahedral group P_{48} or the binary icosahedral group P_{120} with presentations

$$P_{24} = \langle x, y : x^2 = (xy)^3 = y^3, x^4 = 1 \rangle,$$

$$P_{48} = \langle x, y : x^2 = (xy)^3 = y^4, x^4 = 1 \rangle,$$

$$P_{120} = \langle x, y : x^2 = (xy)^3 = y^5, x^4 = 1 \rangle;$$

- The group $P'_{8 \cdot 3^k}$ with presentation

$$P'_{8 \cdot 3^k} = \langle x, y, z : x^2 = (xy)^2 = y^2, zxz^{-1} = y, zyz^{-1} = xy, z^{3^k} = 1 \rangle;$$

- The direct product of any of the above groups with a cyclic group of relatively prime order.

For an analysis of the Seifert structure of the space S^3/G in each of the cases listed above, the reader is referred to Orlik [32].

3.1.4 Surface Bundles

The following notation is used to describe torus and Klein bottle bundles over the circle, as well as spaces formed from pairs of twisted I -bundles $K^2 \tilde{\times} I$.

- $T^2 \times I/A$ represents a torus bundle over the circle, where A is a unimodular 2×2 matrix indicating the homeomorphism under which the torus $T^2 \times 0$ is identified with the torus $T^2 \times 1$. Note that this space is orientable or non-orientable according to whether the determinant of A is 1 or -1 .

More specifically, let μ_0 and λ_0 be closed curves that together generate the fundamental group of the first torus and let μ_1 and λ_1 be the corresponding curves on the second torus. If

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix},$$

then the homeomorphism under which the two tori are identified maps curve μ_0 to $\mu_1^a \lambda_1^c$ and curve λ_0 to $\mu_1^b \lambda_1^d$.

- $K^2 \times I/A$ represents a Klein bottle bundle over the circle, where A is again a unimodular 2×2 matrix indicating the homeomorphism under which the Klein bottle $K^2 \times 0$ is identified with the Klein bottle $K^2 \times 1$. Note that every such space is non-orientable.

Let μ_0 and λ_0 be orientation-preserving and orientation-reversing closed curves respectively on the first Klein bottle that meet transversely in a single point. It is shown in [11] that every element of the fundamental group of this Klein bottle can be represented as $\mu^p \lambda^q$ for some unique pair of integers p and q .

Let μ_1 and λ_1 be the corresponding curves on the second Klein bottle. If

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix},$$

then the homeomorphism under which the two Klein bottles are identified maps curve μ_0 to $\mu_1^a \lambda_1^c$ and curve λ_0 to $\mu_1^b \lambda_1^d$. It is shown in [11] that every such matrix A must be of the form

$$A = \begin{bmatrix} \pm 1 & b \\ 0 & \pm 1 \end{bmatrix}.$$

- $K^2 \tilde{\times} I \cup K^2 \tilde{\times} I/A$ represents a pair of orientable twisted I -bundles $K^2 \tilde{\times} I$ whose corresponding boundary tori are identified, where A is a unimodular 2×2 matrix indicating the homeomorphism under which these boundary tori are identified.

Specifically, let μ_0 and μ_1 be closed curves on the first and second boundary tori each of which projects to a non-trivial orientation-preserving curve on the corresponding core Klein bottle. Let λ_0 and λ_1 be closed curves on the first and second boundary tori each of which represents a double cover of an orientation-reversing curve on the corresponding core Klein bottle. Choose these curves in such a way that together they generate the fundamental groups of the corresponding boundary tori. If

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix},$$

then the homeomorphism under which the two tori are identified maps curve μ_0 to $a\mu_1 + c\lambda_1$ and curve λ_0 to $b\mu_1 + d\lambda_1$.

Once more this notation is consistent with the notation used by Matveev in [29] for orientable 3-manifolds.

3.2 Common Triangulation Components

In order to make the large number of census triangulations easier to both visualise and analyse, we decompose these triangulation into a variety of building blocks. Ideally such building blocks should be large enough that they significantly simplify the representation and analysis of the triangulations containing them, yet small enough that they are frequently reused throughout the different triangulations in each census.

This idea of describing triangulations using medium-sized building blocks has been seen before. Matveev describes in [29] a few such building blocks and in [24] Martelli and Petronio describe a more numerous set of smaller building blocks for orientable triangulations called *bricks*.

As described in the opening remarks of this chapter, the categorisation presented here was developed independently of Martelli and Petronio's results; as such it was intended as an extension of Matveev's results and so the larger building blocks of Matveev have been used as an inspiration for what is presented here. Furthermore, the bricks of Martelli and Petronio whilst ideal for their decomposition theorem presented in [25] are too fine-grained for the visualisation and analysis of triangulations presented in this chapter.

Particular attention must be drawn to the square product spaces described in Sections 3.2.3 and 3.2.4. The construction of these three-tetrahedron and six-tetrahedron building blocks is original, and square product spaces appear within almost all non-orientable triangulations described in this chapter as well as within a number of orientable triangulations.

The first of our building blocks has already been seen; this is the layered solid torus, described in Section 1.2 and used in a variety of proofs in Section 2.6.2. Lemma 2.6.10, which mandates the presence of layered solid tori for triangulations with certain types of associated face pairing graph, already suggests that the layered solid torus will appear frequently within our census triangulations.

We proceed then to define a variety of other common building blocks.

3.2.1 Layered Chains

Recall again the descriptions of the layering process and of layered solid tori presented in Section 1.2. Whereas a layered solid torus is a chain of tetrahedra layered one upon another with a cap at one end (this cap being the Möbius band upon which the first layering takes place), a layered chain is a chain of tetrahedra layered upon each other with nothing to cap it off at either end.

The definition of a layered chain comes with a wealth of associated definitions (upper hinge, top faces and so on); these are introduced primarily to make it easier to describe precisely how a layered chain forms a part of a larger triangulation.

Definition 3.2.1 (Layered Chain) A *layered chain of length n* is a triangulation of a solid torus formed from n tetrahedra as follows.

We begin with the annulus illustrated in Figure 3.1, where edges e are identified and edges h_1 and h_2 are boundary. If we thicken this annulus slightly we can imagine it as a solid torus with four boundary faces (two on each side of the annulus).




Figure 3.1: A two-triangle annulus

The two faces in Figure 3.1 facing out of the page are called the *top faces*, and the two faces facing into the page are called the *bottom faces*. Edges e and f as they appear on the side of the

annulus facing out of the page are called the *outer top edge* and the *inner top edge* respectively. Edges e and f as they appear on the side of the annulus facing into the page are called the *inner bottom edge* and the *outer bottom edge* respectively.

Edges h_1 and h_2 each form a longitude of this solid torus. These are called the *hinge edges* and will remain on the boundary of the layered chain for all lengths n . The hinge edges are oriented so that h_1 runs from edges f to e (the inner top to the outer top edge) and h_2 runs from edges e to f (the inner bottom to the outer bottom edge). Edges h_1 and h_2 are called the *upper hinge* and the *lower hinge* respectively.

A layered chain of length 0 then is simply this annulus. To form a layered chain of length n , we begin with a layered chain of length $n - 1$ and layer a new tetrahedron upon the inner top edge. The old inner top edge becomes internal to the triangulation, the old outer top edge becomes the new inner top edge and the new boundary edge provided by the extra tetrahedron becomes the new outer top edge. The two new boundary faces provided by the extra tetrahedron become the new top faces; note that the old top faces become internal to the triangulation. The hinges, bottom faces and bottom edges remain unchanged.

A layered chain of length 1 is thus illustrated in Figure 3.2. Edge e now forms the inner top edge and edge g is the new outer top edge. Both layered chains of length 0 and 1 are *degenerate* since edge e remains pinched in both triangulations.




Figure 3.2: A layered chain of length 1

Figure 3.3 illustrates a layered chain of length 2, where the dangling face bordered by edges e , g and h_1 on the far right is identified with the face bordered by edges e , g and h_1 on the far left. The inner and outer top edges are now g and j ; the inner and outer bottom edges remain e and f . To obtain a layered chain of length 3, we then layer upon edge g , and so on.




Figure 3.3: A layered chain of length 2

Note that layered chains of length ≥ 2 are no longer degenerate since they no longer contain pinched edges; instead they form triangulations of a solid torus in their own right.

Lemma 3.2.2 In a layered chain of length n , let edges h_1 , h_2 , t_i , t_o , b_i and b_o be the upper hinge, lower hinge, top inner edge, top outer edge, bottom inner edge and bottom outer edge respectively,

as illustrated in Figure 3.4. Orient edges t_i , t_o , b_i and b_o to point from the upper hinge to the lower hinge, and orient the two hinge edges as described in Definition 3.2.1.




Figure 3.4: Orienting the edges of a layered chain

Then the following identities hold, where $+$ represents concatenation of curves and $=$ represents homotopy equivalence.

$$\begin{aligned} t_i &= h_1 + t_o; & t_i &= t_o + h_2; \\ b_o &= h_1 + b_i; & b_o &= b_i + h_2. \end{aligned}$$

Furthermore, hinge edges h_1 and h_2 each individually form longitudinal curves on the boundary torus and the curve $b_o - t_o - (n+1)h_1$ forms a meridinal curve on the boundary torus, where $-$ represents concatenation of a curve in reverse.

Proof The four identities are easily verified by induction on n ; they hold when $n = 0$, and if they hold for $n = k$ it is simple to see from Figure 3.4 that after layering upon edge t_i and relabelling the edges as described in Definition 3.2.1 they will continue to hold for $n = k + 1$.

Similarly, edges h_1 and h_2 are each longitudinal curves when $n = 0$, and layering upon edge t_i does not alter this fact; thus they are longitudinal curves for all n .

Proving the meridinal curve requires a little more work; again we will use induction. Looking back to the layered chain of length 0 in Figure 3.1, we see that running forward along edge f under the bottom of the annulus and then running backward along edge f over the top of the annulus (i.e., running along curve $b_o - t_i$) forms a meridinal curve. So, from the identities above we see that a meridinal curve is $b_o - t_i = b_o - t_o - h_1$ and thus our proposition is true for $n = 0$.

Suppose our proposition is true for $n = k$. When $n = k + 1$ the inner top and outer top edges will have changed; let t_i^- and t_o^- be the inner and outer top edges when $n = k$ and let t_i^+ and t_o^+ be the inner and outer top edges when $n = k + 1$ (so from Definition 3.2.1 we have $t_o^- = t_i^+$). When $n = k + 1$ edges b_o , t_o^- and h_1 are still on the boundary, so from the inductive hypothesis we know $b_o - t_o^- - (k-1)h_1$ is a meridinal curve. From the above identities we see this curve is in fact $b_o - t_i^+ - (k-1)h_1 = b_o - t_o^+ - kh_1$, and so our proposition is true for $n = k + 1$ (and thus for all n). ■

3.2.2 The Triangular Solid Torus

In contrast to the infinite family of building blocks described in Section 3.2.1, we present here a single three-tetrahedron triangulation of a solid torus that is frequently encountered within the orientable census. Note that when phrased in the language of special spines of 3-manifolds, the triangular solid torus is equivalent to the single brick B_4 described by Martelli and Petronio [24].

Definition 3.2.3 (Triangular Solid Torus) A *triangular solid torus* is the specific three-tetrahedron triangulation of a solid torus depicted in Figure 3.5, formed by identifying the two ends of a three-tetrahedron solid triangular prism. This triangulation has six boundary faces which together form a torus.




Figure 3.5: A triangular solid torus

Label the three tetrahedra Δ_1 , Δ_2 and Δ_3 as illustrated. The three edges a_1 , a_2 and a_3 are called the *axis edges*. Observe that each axis edge a_i belongs to tetrahedron Δ_i but neither of the other two tetrahedra, and that the three axis edges individually form three equivalent longitudes of the boundary torus.

For the purpose of identifying directed simple closed curves on the boundary torus, we use any of the directed axis edges to define our longitudinal curve and we use the directed curve $m_1 - n_2 + m_3$ (which bounds the top and bottom faces of the prism) to define our meridinal curve.

Edges m_1 , m_2 and m_3 are called the *major edges*. Observe that each major edge m_i belongs to all three tetrahedra and in particular lies opposite edge a_i in tetrahedron Δ_i . When concatenated in the order $m_1 + m_2 + m_3$ it can be seen that these three edges together form a $(1, 1)$ curve on the boundary torus. The major edges are marked with solid arrowheads in Figure 3.5.

Finally we call edges n_1 , n_2 and n_3 the *minor edges*. Each minor edge n_i can be seen to belong to two of the three tetrahedra and in particular does not belong to tetrahedron Δ_i . When concatenated in the order $n_3 + n_2 + n_1$ we see that the three minor edges together form a $(2, -1)$ curve on the boundary torus. The minor edges are marked with hollow arrowheads in Figure 3.5.

Recall from Definition 2.1.15 that an automorphism is a relabelling that produces precisely the same labelled triangulation. It can be observed that for each permutation p on the set $\{1, 2, 3\}$ there is an automorphism ρ of this triangular solid torus that maps tetrahedron Δ_i to $\Delta_{\rho(i)}$, axis edge a_i to $a_{\rho(i)}$, major edge m_i to $m_{\rho(i)}$ and minor edge n_i to $n_{\rho(i)}$. If p is an odd permutation then the orientation of each edge is reversed under this mapping.

3.2.3 Square Twisted Product Spaces

We now examine a collection of similarly structured three-tetrahedron triangulations of twisted I -bundles over the torus and Klein bottle. Before presenting them individually, Definition 3.2.4 describes the structure common to each of these triangulations.

Definition 3.2.4 (Square Twisted Product Space) A *square twisted product space* is a three-tetrahedron triangulation of either the non-orientable twisted product $T^2 \times I$, the orientable twisted product $K^2 \times I$ or the non-orientable twisted product $K^2 \times I$ formed in the following way.

Let F be a cell decomposition of some closed surface F . We require that F contains precisely one cell which must be a square, and that F contains precisely one vertex. The three possible decompositions F are illustrated in Figure 3.6; the first surface is a torus and the second and third surfaces are both Klein bottles. We refer to this cell decomposition F of the surface S_F as the *central surface* of the square twisted product space. The triangulation that is ultimately constructed will be of a twisted product $S_F \tilde{\times} I$.




Figure 3.6: The three possible central surfaces of a square twisted product space

We replace the single square of F with a single tetrahedron four of whose edges coincide with the four edges of the square, as illustrated in the first and second stages of Figure 3.7. The resulting triangulation represents the surface F thickened everywhere except for at its edges and vertex.




Figure 3.7: Constructing a square twisted product space

By following the edge identifications of the central surface F , the four faces of this new tetrahedron can be arranged to form a square as illustrated in the third stage of Figure 3.7. We call this the *upper square*. The four edges bounding this upper square correspond to the two edges of the tetrahedron that do not lie in F .

We finish the construction by attaching a square pyramid to the upper square as illustrated in the fourth stage of Figure 3.7, so that the surface F can be thickened everywhere. Note that this square pyramid consists of two tetrahedra, and that it can be inserted in two possible ways (one of which is a 90° rotation of the other). Because the resulting structure represents F thickened everywhere, it forms a triangulation of an I -bundle over the surface S_F . Since the faces of the upper square include tetrahedron faces that lie on both sides of F , we see that F forms a one-sided




Figure 3.8: Constructing a square $T^2 \times I$ or a square $K^2 \times I$

surface within this overall triangulation and thus our underlying 3-manifold is in fact a twisted product $S_F \times I$.

The boundary of this triangulation consists of precisely two faces arranged in a square as illustrated in the fifth stage of Figure 3.7. We call this the *boundary square*. Note that the edges of this boundary square are identified to form a double cover of the surface S_F .

Having presented the general construction of a square twisted product space, we proceed to examine the four possible triangulations that can be created using this construction.

Definitions 3.2.5 (Square $T^2 \times I$ and $K^2 \times I$) We define a *square $T^2 \times I$* , a *square non-orientable $K^2 \times I$ of type I and type II* and a *square orientable $K^2 \times I$* to be the four specific triangulations whose constructions are illustrated in Figure 3.8. Each of these triangulations is a square twisted product space as described in Definition 3.2.4 and contains precisely three tetrahedra.

The left diagram of the appropriate portion of Figure 3.8 presents the central surface of the triangulation. The right hand diagram presents the boundary square and its relationship to the central surface. These two diagrams represent the first and last stages of the construction illustrated in Figure 3.7. Note in particular that the triangulation of the boundary square is determined purely by the rotation of the square pyramid that is inserted into the triangulation.

Some basic properties of these triangulations are listed in Table 3.1. In Figures 3.9, 3.10, 3.11 and 3.12 we illustrate each of these triangulations in full detail. Within each figure the two boundary faces are shaded, and in addition the right hand diagram of each of figure illustrates a one-sided torus or Klein bottle at the core of the twisted product $T^2 \times I$ or $K^2 \times I$ respectively. In the left hand diagram of each figure the individual edges are labelled; edges i_1 and i_2 are internal whilst edges m_1 , m_2 and n lie on the boundary. In particular edge n represents the diagonal of the boundary square.

Lemma 3.2.6 Every square twisted product space as described by Definition 3.2.4 is isomorphic to one of the specific triangulations detailed in Definitions 3.2.5.

Triangulation	Orientable	Boundary
Square $T^2 \tilde{\times} I$	No	T^2
Square non-orientable $K^2 \tilde{\times} I$ of type I	No	K^2
Square non-orientable $K^2 \tilde{\times} I$ of type II	No	K^2
Square orientable $K^2 \tilde{\times} I$	Yes	T^2

Table 3.1: Properties of the individual square twisted product spaces




Figure 3.9: A square $T^2 \tilde{\times} I$ and its core torus

Proof This can be proven simply by enumerating all possible square twisted product spaces. In the construction described in Definition 3.2.4 there are three choices for the central surface and two choices for the rotation of the square pyramid, leading to at most six possible triangulations. It is easily established that each of these triangulations is isomorphic to one of the triangulations listed in Definitions 3.2.5. ■

3.2.4 Square Untwisted Product Spaces

In Section 3.2.3 we construct a variety of three-tetrahedron triangulations of twisted I -bundles over the torus and Klein bottle. We can use a modification of this construction to create a collection of six-tetrahedron triangulations of untwisted I -bundles over the torus and Klein bottle as described below. Once more we begin by describing the structure common to each of the triangulations in




Figure 3.10: A square non-orientable $K^2 \tilde{\times} I$ of type I and its core Klein bottle




Figure 3.11: A square non-orientable $K^2 \times I$ of type II and its core Klein bottle




Figure 3.12: A square orientable $K^2 \times I$ and its core Klein bottle

this collection.

Definition 3.2.7 (Square Untwisted Product Space) A *square untwisted product space* is a six-tetrahedron triangulation of either the product $T^2 \times I$ or the product $K^2 \times I$ formed as follows.

Let F be a decomposition into squares of some closed surface S_F . We require that F contains precisely two squares and precisely two vertices, and that each square meets each vertex twice at opposite corners. The three possible decompositions F are illustrated in Figure 3.13; the first surface is a torus and the second and third surfaces are both Klein bottles. We once again refer to this square-based decomposition F of the surface S_F as the *central surface* of the square untwisted product space. The triangulation that is ultimately constructed will be of the product $S_F \times I$.




Figure 3.13: The three possible central surfaces of a square untwisted product space

Label one of the vertices of F the *upper vertex* and the other the *lower vertex*. In Figure 3.13 the upper vertex is coloured black and the lower vertex is coloured white. We then replace each square of F with a tetrahedron four of whose edges coincide with the four edges of the square.

Each tetrahedron additionally has a fifth edge running above the square joining the two instances of the upper vertex as well as a sixth edge running below the square joining the two instances of the lower vertex. These fifth and sixth edges are called the *upper edge* and the *lower edge* respectively of the tetrahedron. This procedure is illustrated in the first and second stages of Figure 3.14.




Figure 3.14: Constructing a square untwisted product space

The resulting triangulation simply represents the central surface F thickened in some places. Note that this figure has precisely eight boundary faces. Four of these boundary faces are adjacent to the two upper edges; by following the edge identifications of the central surface F , these four faces can be arranged to form a large square which we call the *upper square*. Similarly the remaining four boundary faces are adjacent to the lower edges and can be arranged to form another large square which we call the *lower square*. The upper and lower squares are illustrated in the third stage of Figure 3.14. Note that the upper square contains the lower vertex at its centre and the lower square contains the upper vertex at its centre.

Although F has been thickened in some places it has not been thickened everywhere. Specifically the triangulation is pinched at the vertices and edges of F since these vertices and edges belong to both the upper and lower squares. Our final step then is to attach one square pyramid to the upper square and another square pyramid to the lower square as illustrated in the fourth stage of Figure 3.14. Note that each square pyramid is created from two tetrahedra and that each pyramid can be inserted in two possible ways (one of which is a 90° rotation of the other). Since the central surface F is two-sided in the resulting triangulation we see that the underlying 3-manifold is the untwisted I -bundle $S_F \times I$.

Our final triangulation thus uses precisely six tetrahedra. It has four boundary faces arranged into two squares as illustrated in the fifth stage of Figure 3.14 and again in Figure 3.15. The square bounded by the upper edges is called the *upper boundary square* and the square bounded by the lower edges is called the *lower boundary square*. Note that the edges of the upper and lower boundary squares are identified so that each square forms its own copy of the surface S_F .




Figure 3.15: The boundary components of a square untwisted product space

Once more, having outlined the general construction of a square untwisted product space in Definition 3.2.7, we proceed to examine the individual triangulations that can be created in this fashion.

Definitions 3.2.8 (Square $T^2 \times I$) We define a *square $T^2 \times I$ of type I and type II* to be the two specific triangulations of the product space $T^2 \times I$ whose constructions are illustrated in Figure 3.16. Each of these triangulations is a square untwisted product space as described in Definition 3.2.7 and contains precisely six tetrahedra.

The central surface of the triangulation is the torus illustrated in the left hand diagram of the appropriate portion of Figure 3.16. The upper vertex of this central surface is coloured black and the lower vertex is coloured white. The corresponding right hand diagram illustrates the two boundary squares and their relationships to each other and to the central surface. These two diagrams represent the first and last stages of the construction illustrated in Figure 3.14. Note once more that the triangulations of the boundary squares are determined purely by the rotations of the square pyramids that are inserted into the triangulation.




Figure 3.16: Constructing a square $T^2 \times I$ of type I or II

Note that the only way in which the two types of square $T^2 \times I$ differ is in the rotation of the upper square pyramid, which causes the two upper boundary squares to have different diagonal edges.

Definition 3.2.9 (Square $K^2 \times I$) We define a *square $K^2 \times I$ of types I, II, III and IV* to be the four specific triangulations of the product space $K^2 \times I$ whose constructions are illustrated in Figure 3.17. Each of these triangulations is again a square untwisted product space as described in Definition 3.2.7 and contains precisely six tetrahedra.

The central surface of the triangulation is the Klein bottle illustrated in the left hand diagram of the appropriate portion of Figure 3.17, where the upper and lower vertices are again coloured

black and white respectively. Once more the corresponding right hand diagram illustrates the two boundary squares and their relationships to each other and to the central surface.




Figure 3.17: Constructing a square $K^2 \times I$ of type I, II, III or IV

As with the square $T^2 \times I$ triangulations, the only difference between the square $K^2 \times I$ triangulations of types II, III and IV lies in the rotation of the square pyramids. Furthermore, it can be seen with some experimentation that rotating either of the square pyramids in a square $K^2 \times I$ of type I results in an isomorphic triangulation (i.e., another square $K^2 \times I$ of type I). The following result extends this observation to all square untwisted product spaces.

Lemma 3.2.10 Every square untwisted product space as described by Definition 3.2.7 is isomorphic to one of the square $T^2 \times I$ triangulations described in Definitions 3.2.8 or one of the square $K^2 \times I$ triangulations described in Definitions 3.2.9.

Proof As with Lemma 3.2.6, this is easily proven by enumerating all possible square untwisted product spaces. Examining the construction described in Definition 3.2.7 we see that there are three choices for the central surface, two choices for the upper and lower vertices and two choices for the rotation of each square pyramid, leading to at most 24 different square untwisted product spaces. Many of these are trivially seen to be isomorphic to each other and it can be established either by hand or by computer that the remaining triangulations are isomorphic to one of the square $T^2 \times I$ or $K^2 \times I$ triangulations defined earlier.

In the initial derivation of this result a mixture of hand and computer techniques were used, although all results were later verified by hand. Computations were performed using *Regina* which is capable of testing pairs of triangulations for isomorphism. ■

3.3 Families of Closed Triangulations

Having developed a set of medium-sized building blocks in Section 3.2, we can begin combining these to form closed triangulations. In this section we present a number of infinite parameterised families of closed triangulations whose underlying 3-manifolds are known. It can be seen in Sections 3.4 and 3.5 that these families together encompass almost all prime minimal triangulations containing up to six tetrahedra.

As discussed in the opening remarks of this chapter, such a collection of parameterised families of 3-manifold triangulations is a useful tool. When analysing 3-manifolds and their triangulations, such families enable us to study large classes of triangulations simultaneously. In addition, because the families presented here are infinite, we obtain a means for recognising the underlying 3-manifolds of many new triangulations that we are encountering for the first time.

Matveev begins this process of constructing families of orientable triangulations in [29]. Although his results are presented in the language of special spines of 3-manifolds, he constructs families of spines analogous to layered lens spaces, twisted layered loops and augmented triangular solid tori, discussed in Sections 3.3.1, 3.3.2 and 3.3.4 respectively.

Martelli and Petronio continue this construction of orientable families in [24]. Like Matveev they work with special spines of 3-manifolds, constructing families of spines that correspond to layered chain pairs and certain restricted types of plugged triangular solid torus, discussed in Sections 3.3.3 and 3.3.6 respectively. The constructions described in this chapter were developed independently of the work of Martelli and Petronio.

In the following sections we describe the constructions listed above and in some cases generalisations of these constructions, all presented within the context of 3-manifold triangulations. In addition to these constructions we present new families of triangulations including chained triangular solid tori, generalised plugged triangular solid tori, square torus and Klein bottle bundles and square product pairs, discussed in Sections 3.3.5, 3.3.6, 3.3.7 and 3.3.8 respectively. The square torus and Klein bottle bundles are of particular interest since these families include non-orientable triangulations.

3.3.1 Layered Lens Spaces

It is known that all lens spaces can be constructed from two solid tori by identifying their torus boundaries according to some homeomorphism. It is expected then that we can create triangulations of lens spaces by constructing two layered solid tori and identifying the corresponding boundary faces. The resulting triangulations are called layered lens spaces and are described more formally below.

The layered lens space construction is well known. Jaco and Rubinstein present this construction in [17], and an analogous construction involving special spines of 3-manifolds is presented by Matveev in [29].

Definition 3.3.1 (Layered Lens Space) Let L_1 and L_2 be two layered solid tori at least one of which is non-degenerate, as described by Definition 1.2.2. We construct a new triangulation by identifying the two boundary faces of L_1 with the two boundary faces of L_2 in such a way that

these identifications represent a homeomorphism between the two boundary tori. This procedure is illustrated in Figure 3.18. The resulting triangulation is called a *layered lens space*.




Figure 3.18: Constructing a layered lens space

Lemma 3.3.2 Every layered lens space is indeed a triangulation of some lens space, where we consider S^3 and $S^2 \times S^1$ to be the lens spaces $L(1, 0)$ and $L(0, 1)$ respectively.

Proof This is clear from Definition 3.3.1, since every space constructed from two solid tori by identifying their torus boundaries according to some homeomorphism is indeed a lens space. ■

Lemma 3.3.3 When constructing a layered lens space from layered solid tori L_1 and L_2 , the identification of the boundaries of L_1 and L_2 in fact represents a layering of the outermost tetrahedron of L_1 onto a boundary edge of L_2 (and vice versa).

Proof This is easily verified by examining the twelve possible boundary identifications that correspond to homeomorphisms of the boundary tori. ■

Lemma 3.3.4 Consider a layered lens space constructed from layered solid tori L_1 and L_2 , where L_2 contains at least one tetrahedron. Let T be the boundary tetrahedron of L_2 . Then both $L_1 \cup \{T\}$ and $L_2 \setminus \{T\}$ are also layered solid tori, and this same layered lens space can be constructed as described in Definition 3.3.1 from the layered solid tori $L_1 \cup \{T\}$ and $L_2 \setminus \{T\}$.

Proof $L_1 \cup \{T\}$ is a layered solid torus by Lemma 3.3.3. $L_2 \setminus \{T\}$ is a layered solid torus since Definition 1.2.2 requires the boundary tetrahedron of the layered solid torus L_2 to be layered onto the boundary of some smaller layered solid torus. It can then be established by following face and edge identifications that the layering of tetrahedron T onto the boundary of $L_2 \setminus \{T\}$ corresponds to a homeomorphism between the torus boundaries of $L_1 \cup \{T\}$ and $L_2 \setminus \{T\}$. ■

Lemma 3.3.4 effectively allows us to move tetrahedra back and forth between the two layered solid tori described in Definition 3.3.1 whilst preserving the underlying structure of a layered lens space.

Corollary 3.3.5 Any layered lens space formed from n tetrahedra can be constructed as described in Definition 3.3.1 by identifying the boundaries of an n -tetrahedron layered solid torus and a 0-tetrahedron layered solid torus, i.e., a Möbius band.

Proof Let T be a layered lens space constructed from layered solid tori L_1 and L_2 . We simply use Lemma 3.3.4 to repeatedly move tetrahedra across from L_2 to L_1 until L_1 contains n tetrahedra and L_2 contains none. ■

Theorem 3.3.6 Let T be a layered lens space constructed by identifying the boundaries of an $\text{LST}(p, q, r)$ and a 0-tetrahedron $\text{LST}(1, 1, -2)$, i.e., a Möbius band. Furthermore, let these boundaries be identified so that the edge with parameter r of the $\text{LST}(p, q, r)$ is identified with the edge of the $\text{LST}(1, 1, -2)$ with parameter -2 , i.e., the boundary edge of the Möbius band. Then T is a triangulation of the lens space $L(|p - q|, |p|)$.

Proof Let R be the two-triangle torus that forms the common boundary of the two layered solid tori described above. Recalling the description of a three-parameter torus curve in Definition 1.2.3, we can label the edges and faces of R in such a way that the meridinal curve of the $\text{LST}(p, q, r)$ forms a (p, q, r) curve on the torus R and the meridinal curve of the Möbius band forms a $(1, 1, -2)$ curve on the same torus R . To identify the lens space that has been constructed, we must express the meridinal curve of the Möbius band as a combination of longitudes and meridians of the $\text{LST}(p, q, r)$.

Select some integers x and y for which $py - qx = 1$. This is possible since p and q are coprime by Lemma 1.2.15. Note that the resulting integers x and y must also be coprime, since any common factor of x and y must divide $py - qx = 1$. From Theorem 1.2.19 it follows that the standard representations of a (p, q, r) curve and an $(x, y, -x - y)$ curve intersect in precisely $py - qx = 1$ point. Thus the $(x, y, -x - y)$ curve on the torus R forms a longitude of the $\text{LST}(p, q, r)$.

Examine now the torus curve that is constructed by combining $(y - x)$ copies of a (p, q, r) curve with $(p - q)$ copies of a $(x, y, -x - y)$ curve. The resulting curve, by Lemma 1.2.6, is a

$$\begin{aligned} & ((y - x)p + (p - q)x, (y - x)q + (p - q)y, (y - x)(-p - q) + (p - q)(-x - y)) \\ = & (py - qx, py - qx, -2(py - qx)) \\ = & (1, 1, -2) \end{aligned}$$

curve on the torus. Thus the meridinal curve of the Möbius band is identified with a curve on the $\text{LST}(p, q, r)$ formed from $(p - q)$ longitudes and $(y - x)$ meridians. We see then that our triangulation must be of the lens space $L(|p - q|, |y - x|)$.

Finally, we observe that $p(y - x) = (py - qx) - (p - q)x = 1 - (p - q)x$, and so $|p| \cdot |y - x| \equiv \pm 1 \pmod{|p - q|}$. Thus the lens space constructed can also be represented as $L(|p - q|, |p|)$. ■

Algorithm 3.3.7 (Layered Lens Space Construction) From Theorem 3.3.6 we see that we can construct the lens space $L(p, q)$ for any coprime integers $p, q \geq 0$ as follows. We construct a non-degenerate $\text{LST}(p - 2q, q, q - p)$ as described in Algorithm 1.2.17. We then join this layered solid torus to a Möbius band so that the boundary edge of the Möbius band is identified with the edge whose parameter is $p - 2q$.

As a special case, in the interests of minimising the number of tetrahedra in the triangulation, we construct the 3-sphere $L(1, 0)$ differently as follows. We construct an $\text{LST}(-3, 2, 1)$ from a single tetrahedron, and then join this layered solid torus to a Möbius band so that the boundary edge of the Möbius band is identified with the edge whose parameter is -3 .

The layered lens space thus constructed is denoted by the symbol $L_{p,q}$. Note that $L(p,q)$ denotes the underlying 3-manifold whereas $L_{p,q}$ denotes this specific triangulation of $L(p,q)$.

When constructing the lens space $L(p,q)$, the number of tetrahedra in the resulting triangulation can be reduced by ensuring that parameter q is selected so that $q \leq \frac{1}{2}p$. This is possible using the result that lens spaces $L(p,q)$ and $L(p, kp \pm q)$ are homeomorphic for any $k \in \mathbb{Z}$.

3.3.2 Layered Loops

A layered loop is simply a layered chain whose top faces are identified with its bottom faces, causing the chain to loop back upon itself. This identification of faces can be performed with or without a twist, resulting in the two types of layered loop described below.

In the case of a twisted layered loop as described in Definition 3.3.9, Matveev presents in [29] an analogous construction involving special spines of 3-manifolds.

Definition 3.3.8 (Untwisted Layered Loop) Consider a layered chain of length $n \geq 1$ as described by Definition 3.2.1 and illustrated in Figure 3.19. Let h_1 , h_2 , t_i , t_o , b_i and b_o denote the upper hinge, lower hinge, top inner edge, top outer edge, bottom inner edge and bottom outer edge respectively.




Figure 3.19: A layered chain of length n

An *untwisted layered loop of length n* , denoted by the symbol C_n , is formed by layering the top tetrahedron upon the inner bottom edge, or equivalently by layering the bottom tetrahedron upon the inner top edge. This layering is performed in such a way that upper hinge h_1 is identified with itself and lower hinge h_2 is identified with itself. In particular, face ACB is identified with face WZX and face CBD is identified with face YWZ . Note that as a result edges t_i and b_o are identified and edges t_o and b_i are identified. The final edge identifications are illustrated in Figure 3.20.

Definition 3.3.9 (Twisted Layered Loop) A *twisted layered loop of length $n \geq 1$* , denoted by the symbol \tilde{C}_n , is constructed identically to an untwisted layered loop of length n with the exception that the top faces are rotated by 180° before being identified with the bottom faces.

Specifically, we begin again with the layered chain of length n depicted in Figure 3.19. Once more we layer the top tetrahedron upon the inner bottom edge (or equivalently layer the bottom tetrahedron upon the inner top edge), but this time we layer in such a way that the upper and lower hinges h_1 and h_2 are identified. In particular, faces ACB and ZWY are identified and faces




Figure 3.20: An untwisted layered loop of length n

CBD and XZW are identified. Note again that as a result edges t_i and b_o are identified and edges t_o and b_i are identified. The final edge identifications are illustrated in Figure 3.21.




Figure 3.21: A twisted layered loop of length n

Although the layered chain of length 1 is degenerate, it is worth noting that both the untwisted and twisted layered loops of length 1 are full triangulations in their own right. This is because the link of the pinched edge in the degenerate layered chain is fleshed out to become a full circle when the top and bottom faces of the layered chain are identified.

Theorem 3.3.10 For each $n \geq 1$, the untwisted layered loop C_n is a two-vertex triangulation of the lens space $L(n, 1)$.

Proof Within each tetrahedron of the layered loop, insert a quadrilateral between the upper and lower hinges. Together these quadrilaterals form a torus R within the layered loop as illustrated in the left hand diagram of Figure 3.22. It can be seen that this torus divides the underlying 3-manifold into two solid tori, each with a hinge edge at its core. It follows that the torus R represents a genus one Heegaard splitting of the underlying 3-manifold.




Figure 3.22: A two-sided torus within an untwisted layered loop

The torus R is illustrated again in the right hand diagram of Figure 3.22 for the case $n = 4$. The curve m_1 represents a meridinal curve on the boundary of the upper solid torus and the curve m_2 represents a meridinal curve on the boundary of the lower solid torus, where in the general case the curve m_1 is formed from precisely n arcs across the diagram. It is clear from this diagram that the Heegaard splitting described by the torus R is in fact a Heegaard splitting of the lens space $L(n, 1)$. ■

Theorem 3.3.11 For each $n \geq 1$, the twisted layered loop \tilde{C}_n is a one-vertex triangulation of the space S^3/Q_{4n} , or equivalently of the Seifert fibred space $\text{SFS}(S^2 : (2, 1), (2, 1), (n, -n + 1))$.

Proof Once more we insert a quadrilateral into each tetrahedron between the upper and lower hinges. In the case of a twisted layered loop, these quadrilaterals combine to form a one-sided Klein bottle K as illustrated in the left hand diagram of Figure 3.23. Slicing along this Klein bottle reduces the underlying 3-manifold to a single solid torus with the hinge edge h at its core. This Klein bottle thus represents a one-sided Heegaard splitting of non-orientable genus two of the underlying 3-manifold.




Figure 3.23: A one-sided Klein bottle within a twisted layered loop

The Klein bottle K is illustrated once more in the right hand diagram of Figure 3.23 for the case $n = 4$. The curve m represents a meridinal curve on the boundary of the solid torus described above; note that in the general case this curve consists of a single long arc crossing the diagram from top to bottom and precisely n short arcs crossing the diagram from right to left.

Rubinstein in [38] investigates 3-manifolds with one-sided Heegaard splittings of non-orientable genus two, and in particular presents a method for identifying such 3-manifolds based upon a meridinal curve of the corresponding solid torus. Using the meridinal curve illustrated in Figure 3.23, the results of Rubinstein show the underlying 3-manifold to be the space S^3/Q_{4n} . Orlik then shows in [32] that the space S^3/Q_{4n} can in fact be represented as the Seifert fibred space $\text{SFS}(S^2 : (2, 1), (2, 1), (n, -n + 1))$. ■

3.3.3 Layered Chain Pairs

In Section 3.3.2 we form a layered loop by constructing a single layered chain and identifying its two ends. We now investigate a style of triangulation in which we construct two layered chains and identify faces from both ends of the first chain with faces from both ends of the second.

An analogous constructing involving special spines of 3-manifolds is presented by Martelli and Petronio in [24]. In particular the layered chain pair $C_{r,s}$ as described by Definition 3.3.12 corresponds to the closed brick $C_{r,s}$ as presented in [24].

Definition 3.3.12 (Layered Chain Pair) For each pair of positive integers r and s we construct an (r, s) *layered chain pair*, denoted by the symbol $C_{r,s}$, as follows.

Let T be a layered chain of length r and let T' be a layered chain of length s . Let h_1, h_2, t_i, t_o, b_i and b_o denote the upper hinge, lower hinge, top inner edge, top outer edge, bottom inner edge and bottom outer edge respectively of T . Let $h'_1, h'_2, t'_i, t'_o, b'_i$ and b'_o denote the upper hinge, lower hinge, top inner edge, top outer edge, bottom inner edge and bottom outer edge respectively of T' .

Together these two layered chains have eight boundary faces. We then identify the four boundary faces of T with the four boundary faces of T' as illustrated in Figure 3.24. In particular, these faces are identified in such a way that the following edge identifications result.

- The two hinge edges of T are identified with the top and bottom inner edges of T' ;
- The top and bottom inner edges of T are identified with the two hinge edges of T' ;
- The top and bottom outer edges of T and the top and bottom outer edges of T' are all identified as a single edge.




Figure 3.24: Constructing a layered chain pair

Lemma 3.3.13 The layered chain pairs $C_{r,s}$ and $C_{s,r}$ are isomorphic triangulations. Furthermore, the layered chain pair $C_{n-1,1}$ is in fact simply the twisted layered loop \tilde{C}_n .

Proof These equivalences can be established through a careful examination and comparison of each of the corresponding constructions. ■

Theorem 3.3.14 For each $r, s \geq 1$, the layered chain pair $C_{r,s}$ is a triangulation of the Seifert fibred space $\text{SFS}(S^2 : (2, -1), (r+1, 1), (s+1, 1))$.

Proof This result can be established by constructing a fibration of the entire triangulation. Such a fibration is too complex to depict here. Nevertheless, the exceptional fibres are easily identified. The layered chains of lengths r and s have exceptional fibres of index r and s respectively at the cores of the solid tori that they represent. Furthermore, the single edge to which the top and bottom outer edges of both layered chains are identified forms an exceptional fibre of index two. ■

3.3.4 Augmented Triangular Solid Tori

We move now to examine a series of constructions based upon the triangular solid torus. The first of these constructions is the augmented solid torus, which involves attaching three layered solid tori to the six boundary faces of a triangular solid torus.

The augmented triangular solid torus is the last of our constructions for which a corresponding construction is described by Matveev in [29] for special spines of 3-manifolds.

Definition 3.3.15 (Augmented Triangular Solid Torus) Let $\alpha_1, \mu_1, \alpha_2, \mu_2, \alpha_3$ and μ_3 be integers for which each α_i is strictly positive and for which α_i and μ_i are relatively prime for each i . An *augmented triangular solid torus* with parameters $(\alpha_1, \mu_1, \alpha_2, \mu_2, \alpha_3, \mu_3)$, denoted by the symbol $A_{\alpha_1, \mu_1 | \alpha_2, \mu_2 | \alpha_3, \mu_3}$, is the triangulation formed using the following construction.




Figure 3.25: Annuli on the boundary of a triangular solid torus

We begin with a triangular solid torus as illustrated in the left hand diagram of Figure 3.25. Observe that the axis edges a_1, a_2 and a_3 partition the six boundary faces of this triangular solid torus into three annuli as illustrated in the right hand diagram of Figure 3.25. We label these annuli S_1, S_2 and S_3 , where the annulus S_i contains major edge m_i and minor edge n_i . To each annulus S_i we attach an $LST(\alpha_i, \mu_i, -\alpha_i - \mu_i)$, where the layered solid torus $LST(\alpha_i, \mu_i, -\alpha_i - \mu_i)$ is constructed according to Algorithm 1.2.17. This layered solid torus is attached so that the edges with parameters α_i and μ_i are identified with the axis edges and major edges of annulus S_i respectively. Note that as a result of this construction all three axis edges are identified. The attachment of a layered solid torus to annulus S_1 is illustrated in Figure 3.26.

Note that degenerate layered solid tori may be used, and in particular that the two faces of an annulus may simply be identified with each other by attaching the 0-tetrahedron $LST(2, -1, -1)$. For brevity, if a pair α_i, μ_i is omitted from the symbolic name of an augmented triangular solid torus then this pair is assumed to be $2, -1$. For instance, the augmented triangular solid torus $A_{2, -5}$ is in fact $A_{2, -5 | 2, -1 | 2, -1}$.

Lemma 3.3.16 The order of the three pairs α_i, μ_i in the symbolic name of an augmented triangular solid torus is irrelevant. That is, interchanging the pair α_i, μ_i with some other pair α_j, μ_j results in an isomorphic triangulation.




Figure 3.26: Constructing an augmented triangular solid torus

Proof This is straightforward to establish as a result of the symmetries of the triangular solid torus and the layered solid torus. ■

Theorem 3.3.17 The augmented triangular solid torus $A_{\alpha_1, \mu_1 | \alpha_2, \mu_2 | \alpha_3, \mu_3}$ is a triangulation of the Seifert fibred space $SFS(S^2 : (\alpha_1, \mu_1), (\alpha_2, \mu_2), (\alpha_3, \mu_3), (1, 1))$.

Proof As with layered chain pairs, this result can be proven by constructing a fibration of the entire triangulation. The fibres within the triangular solid torus run vertically from top to bottom. An exceptional fibre of index α_i is then formed at the core of the layered solid torus corresponding to the pair α_i, μ_i for each i . ■

3.3.5 Chained Triangular Solid Tori

The chained triangular solid torus is the next in our suite of triangulations based upon the triangular solid torus. This particular construction involves the attachment of one layered solid torus and one layered chain to the six boundary faces of a triangular solid torus.

There are two distinct types of chained triangular solid torus, these being the axial type and the major type. Each is described in turn.

Definition 3.3.18 (Chained Triangular Solid Torus, Axial) Let k, α and μ be integers for which k and α are strictly positive and for which α and μ are relatively prime. A *chained triangular solid torus of axial type* with parameters (k, α, μ) , denoted by the symbol $X_{k|\alpha, \mu}$, is the triangulation formed using the following construction.

We begin with a triangular solid torus T and a layered chain L of length k , as illustrated in Figure 3.27. Observe that the axis edges a_1, a_2 and a_3 partition the six boundary faces of T into three annuli S_1, S_2 and S_3 , where each annulus S_i contains major edge m_i and minor edge n_i . These are the same annuli as described in Definition 3.3.15 and are illustrated in Figure 3.25.

To two of these annuli we attach the two ends of the layered chain L . This procedure is illustrated in Figure 3.27, where the layered chain is attached to annuli S_1 and S_3 . The layered chain must be attached precisely as illustrated in the figure, where h_1, h_2, t_i, t_o, b_i and b_o denote the upper hinge, lower hinge, top inner edge, top outer edge, bottom inner edge and bottom outer edge respectively of L .




Figure 3.27: Constructing a chained triangular solid torus of axial type

Note that the two top faces of L are identified with the two faces of one annulus (S_3 in the diagram) and the two bottom faces of L are identified with the two faces of the other annulus (S_1 in the diagram). Furthermore, the following edge identifications can be observed.

- The outer top and outer bottom edges of L are identified with major edges of T ;
- The inner top and inner bottom edges of L are identified with minor edges of T ;
- The two hinge edges of L and the three axis edges of T are all identified as a single edge.

To the remaining annulus we attach an $\text{LST}(\alpha, \mu, -\alpha - \mu)$ in such a way that the edges with parameters α and μ in the layered solid torus are identified with the axis edges and major edges of the annulus respectively. As in Definition 3.3.15, this layered solid torus may be degenerate and may in fact be the 0-tetrahedron $\text{LST}(2, -1, -1)$.

Definition 3.3.19 (Chained Triangular Solid Torus, Major) Let k , α and μ be integers for which k and α are strictly positive and for which α and μ are relatively prime. A *chained triangular solid torus of major type* with parameters (k, α, μ) , denoted by the symbol $J_{k|\alpha, \mu}$, is the triangulation formed using the following construction.

As with a chained triangular solid torus of axial type, we begin with a triangular solid torus T and a layered chain L of length k as illustrated in Figure 3.28. Once more we partition the six boundary faces of T into three annuli S_1 , S_2 and S_3 as illustrated in Figure 3.25.

To two of these annuli we attach the two ends of the layered chain L . This attachment however is performed differently from the attachment in a chained triangular solid torus of axial type. Each of the two annuli has a face identified with both a top and a bottom face of L . The full identification of faces is illustrated in Figure 3.28, where the layered chain is once more attached




Figure 3.28: Constructing a chained triangular solid torus of major type

to annuli S_1 and S_3 . The layered chain must be attached precisely as illustrated in the diagram. As a result of this attachment of the layered chain and the triangular solid torus, the following edge identifications can be observed.

- The two hinge edges of L are identified with major edges of T ;
- The inner top and inner bottom edges of L are identified with axis edges of T , and two of these axis edges are identified;
- The outer top and outer bottom edges of L and two of the minor edges of T are all identified as a single edge.

Once more we attach an $LST(\alpha, \mu, -\alpha - \mu)$ to the remaining annulus on the boundary of T so that the edges with parameters α and μ in the layered solid torus are identified with the axis edges and major edges of the annulus respectively.

Theorem 3.3.20 The chained triangular solid torus $X_{k|\alpha, \mu}$ is a triangulation of the Seifert fibred space $SFS(S^2 : (2, 1), (2, -1), (k\alpha - \mu, \alpha))$.

Proof We establish this result by constructing a fibration of the entire triangulation. The full fibration is too complex to depict here, but we can nevertheless describe the exceptional fibres. An exceptional fibre of index $(k\alpha - \mu)$ is found at the core of the layered solid torus. Furthermore, using the notation of Figure 3.27, two exceptional fibres each of index two run parallel to the edge t_i through the Klein bottle formed by the top two faces of the layered chain. ■

Theorem 3.3.21 The chained triangular solid torus $J_{k|\alpha, \mu}$ is a triangulation of the Seifert fibred space $SFS(S^2 : (2, 1), (k + 1, 1), (\alpha - \mu, \mu))$.

Proof Once more we can construct a fibration of the entire triangulation. This time we find exceptional fibres of index $\alpha - \mu$ and $(k + 1)$ at the cores of the layered solid torus and the solid torus formed by the layered chain. Using the notation of Figure 3.28, the minor edges n_1 and n_3 which are identified form a third exceptional fibre of index two. ■

3.3.6 Plugged Triangular Solid Tori

Our final construction involving the triangular solid torus is the plugged triangular solid torus. This construction entails creating a triangular solid torus, layering tetrahedra upon the three boundary annuli and then inserting a two-tetrahedron plug to close off the new boundary that results.

In Definition 3.3.22 we define the plugged triangular solid tori P_{k_1, k_2, k_3} and P'_{k_1, k_2, k_3} . In [24] Martelli and Petronio present a construction involving special spines of 3-manifolds that corresponds to the class of triangulations $P_{k, 0, 0}$ for $k \geq 0$. In Martelli and Petronio's notation such a triangulation is analogous to the closed brick E_k .

Definition 3.3.22 (Plugged Triangular Solid Torus) A *plugged triangular solid torus* is a triangulation formed according to the following construction. Such a triangulation may be of either major type or minor type; each of these is discussed in detail at a later point in the definition.

Let T be a triangular solid torus. Once more we consider the six boundary faces of T to be partitioned into annuli S_1 , S_2 and S_3 , where annulus S_i contains major edge m_i and minor edge n_i and is bordered by two axis edges. These annuli are illustrated in Figure 3.25 on page 108.

For each annulus S_i we initially define the *outer major edge* m'_i to be the major edge m_i and the *outer minor edge* n'_i to be the minor edge n_i . These definitions are subject to change as described below.

Recall the layering process described in Definition 1.2.1. Upon each annulus S_i we may layer tetrahedra as follows.

- A *positive layering* is the process of layering a new tetrahedron upon the outer major edge m'_i . We redefine the outer major and minor edges as follows. The old outer major edge becomes internal to the triangulation. The old outer minor edge becomes the new outer major edge. The new diagonal edge provided by the additional tetrahedron becomes the new outer minor edge. Such a layering is depicted in the left hand diagram of Figure 3.29.
- A *negative layering* is the process of layering a new tetrahedron upon the outer minor edge n'_i . Again we redefine the outer major and minor edges as follows. The old outer minor edge becomes internal to the triangulation. The old outer major edge becomes the new outer minor edge. The new diagonal edge provided by the additional tetrahedron becomes the new outer major edge. A negative layering is depicted in the right hand diagram of Figure 3.29.

After some sequence of positive and negative layerings has been performed, the resulting triangulation once more has six boundary faces partitioned into three annuli S'_1 , S'_2 and S'_3 , where the annulus S'_i is the portion of the boundary obtained from the original annulus S_i by layering tetrahedra upon S_i . The final stage of the construction is then to attach a two-tetrahedron plug to these six boundary faces.




Figure 3.29: A positive and a negative layering upon the annulus S_i

The plug may be attached in one of two ways. A plug of *major type* is attached in such a way that the three edges e_1 , e_2 and e_3 that form the equator of the plug are identified with the three outer major edges m'_1 , m'_2 and m'_3 respectively. A plug of *minor type* is attached so that these same edges e_1 , e_2 and e_3 are identified with the three outer minor edges n'_1 , n'_2 and n'_3 respectively.

The processes of attaching a plug of major or minor type are illustrated in Figures 3.30 and 3.31 respectively. In each case the attachment of the plug must be performed precisely as illustrated in the diagram.




Figure 3.30: Attaching a plug of major type

Note that in each of Figures 3.30 and 3.31 the annuli S'_1 , S'_2 and S'_3 are depicted so that the corresponding outer major edges appear horizontal. This is simply for ease of presentation, and in no way implies that the outer major edges together bound a meridinal disc of the solid torus contained within these annuli.

It can be observed that when attaching a plug of major type, each axis edge of the central triangular solid torus is identified with an outer minor edge. Similarly, when attaching a plug of minor type each axis edge of the triangular solid torus is identified with an outer major edge.

The specific parameterised plugged triangular solid tori are then defined as follows.

- The *plugged triangular solid torus of major type* with parameters (k_1, k_2, k_3) , denoted by the symbol P_{k_1, k_2, k_3} , is formed as described above by attaching a plug of major type. The




Figure 3.31: Attaching a plug of minor type

meanings of the parameters k_i are described below.

- The *plugged triangular solid torus of minor type* with parameters (k_1, k_2, k_3) , denoted by the symbol P'_{k_1, k_2, k_3} , is similarly formed by attaching a plug of minor type. Again the meanings of the parameters k_i are described below.

Each parameter k_i in the above triangulations describes the specific layerings that take place upon the annulus S_i . If $k_i = 0$ then no layerings take place. If $k_i > 0$ then a sequence of k_i positive layerings are performed, and if $k_i < 0$ then a sequence of $|k_i|$ negative layerings are performed.

For brevity, if a parameter k_i is omitted from the symbolic name of a plugged triangular solid torus then it is assumed to be zero. For instance, the triangulation P_2 is in fact $P_{2,0,0}$.

Lemma 3.3.23 The order of the three parameters k_1, k_2 and k_3 in the symbolic name of a plugged triangular solid torus is irrelevant. That is, reordering these parameters results in an isomorphic triangulation.

Proof This is a simple consequence of the symmetries of the triangular solid torus, the layered chain and the two-tetrahedron plug. ■

Theorem 3.3.24 The plugged triangular solid torus P_{k_1, k_2, k_3} is a triangulation of the Seifert fibred space $\text{SFS}(S^2 : (2, -1), (3, 1), (5 + k_1 + k_2 + k_3))$. The plugged triangular solid torus P'_{k_1, k_2, k_3} is a triangulation of the Seifert fibred space $\text{SFS}(S^2 : (2, -1), (3, 1), (4 - k_1 - k_2 - k_3))$.

Proof A fibration of each of these plugged triangular solid tori can be constructed which is again too complex to present here in full detail. An exceptional fibre is found running vertically through the centre of the triangular solid torus. This exceptional fibre has index $(5 + k)$ or $(4 - k)$ according to whether the plugged triangular solid torus is of major type or minor type. An exceptional fibre of index three runs horizontally through the centre of the two-tetrahedron plug from left to right, and an exceptional fibre of index two horizontally through the two-tetrahedron plug passing through the midpoints of each of the axis edges. ■

3.3.7 Square Surface Bundles

The remaining families of triangulations presented in this chapter are based not upon the triangular solid torus but upon the square product spaces of Sections 3.2.3 and 3.2.4. These families are of particular interest because they include non-orientable as well as orientable triangulations.

The first of these constructions is the square surface bundle, formed from a square untwisted product space by identifying the top and bottom boundary surfaces under some homeomorphism. A layering of tetrahedra between these surfaces may be used to achieve the desired homeomorphism.

Definition 3.3.25 (Square Surface Bundle) A *square surface bundle* is a triangulation formed using the following construction. Note that square surface bundles come in various types; the specific triangulations $T_{t|p,q|r,s}$ and $K_{t|p,q|r,s}$ are detailed at a later point in the definition.

Let S be a six-tetrahedron square untwisted product space as described in Definition 3.2.7. From Lemma 3.2.10 we see that S must be a square $T^2 \times I$ of type I or II or a square $K^2 \times I$ of type I, II, III or IV. Assume then that the boundary squares of S are arranged as illustrated in Figure 3.16 on page 99 for a square $T^2 \times I$ or as illustrated in Figure 3.17 on page 100 for a square $K^2 \times I$.

Let $\alpha_1, \beta_1, \alpha_2$ and β_2 be the directed edges on the upper and lower boundary squares illustrated in Figure 3.32. Note that for all six square untwisted product spaces, the curves α_1 and β_1 together generate the fundamental group of the upper boundary surface and the curves α_2 and β_2 together generate the fundamental group of the lower boundary surface.




Figure 3.32: Edges $\alpha_1, \beta_1, \alpha_2$ and β_2 on the boundary of a square product space

Let p, q, r and s be integers. We aim to identify the upper and lower boundary surfaces using a homeomorphism that identifies directed edge α_1 with the curve $\alpha_2^p \beta_2^q$ and directed edge β_1 with the curve $\alpha_2^r \beta_2^s$.

For some values of p, q, r and s this identification of surfaces can be realised by an immediate identification of the corresponding boundary faces. On the other hand, for some values of p, q, r and s the curves $\alpha_2^p \beta_2^q, \alpha_2^r \beta_2^s$ and the curve corresponding to the upper boundary diagonal edge are not realised as edges of the lower boundary surface. In this case we must layer tetrahedra upon the lower boundary surface so that two new lower boundary faces are created whose edges form the required curves on the underlying surface. Note that since every triangulation described in this chapter is formed from at most six tetrahedra, none of these triangulations require this additional layering process.

For some values of p, q, r and s it is not possible to identify the upper and boundary surfaces as described, since there is no homeomorphism between the two surfaces that identifies the upper boundary edges with the required curves on the lower boundary surface.

The *square torus bundles* $T_{I|p,q|r,s}$ and $T_{II|p,q|r,s}$ are the specific triangulations formed as described above where the initial product space S is a square $T^2 \times I$ of type I or II respectively. The *square Klein bottle bundles* $K_{I|p,q|r,s}$, $K_{II|p,q|r,s}$, $K_{III|p,q|r,s}$ and $K_{IV|p,q|r,s}$ are the specific triangulations formed as described above where the initial product space S is a square $K^2 \times I$ of type I, II, III or IV respectively.

Example 3.3.26 Consider the triangulation $T_{II|1,-1|0,1}$. The square product space from which this triangulation is constructed is a square $T^2 \times I$ of type II, illustrated in the first two diagrams of Figure 3.33. The homeomorphism under which the upper and lower boundary surfaces are identified must identify directed edges α_1 and β_1 with the curves $\alpha_2\beta_2^{-1}$ and β_2 respectively. Such a homeomorphism can be realised by identifying the upper and lower boundary faces as illustrated in the right hand diagram of Figure 3.33. In particular, faces ABC and WZY are identified and faces BCD and XWZ are identified.




Figure 3.33: Constructing the square torus bundle $T_{II|0,-1|-1,1}$

Theorem 3.3.27 For each set of integers p , q , r and s for which the corresponding triangulations can be constructed, the underlying 3-manifolds of the square surface bundles with parameters p , q , r and s are as follows.

- $T_{I|p,q|r,s}$ and $T_{II|p,q|r,s}$ are both triangulations of the space $T^2 \times I / \begin{bmatrix} p & r \\ q & s \end{bmatrix}$;
- $K_{II|p,q|r,s}$, $K_{III|p,q|r,s}$ and $K_{IV|p,q|r,s}$ each represent the space $K^2 \times I / \begin{bmatrix} \{p\} - \{r\} & \{r\} \\ 0 & s - r \end{bmatrix}$, where the symbol $\{x\}$ is defined to be 1 if x is odd and 0 if x is even and where $p + r$ is odd, $|p - q| = |r - s| = 1$ and $p - q + r - s = 0$.

Proof Each of these results is verified by matching curves on the upper and lower boundary surfaces as described in Section 3.1.4. ■

Identifying the underlying 3-manifold of the triangulation $K_{I|p,q|r,s}$ is more difficult and is not described here. The underlying 3-manifold of each square Klein bottle bundle formed from six tetrahedra is however listed in the tables of Appendix A.

3.3.8 Square Product Pairs

To conclude our series of constructions we present the square product pair, formed from a pair of square twisted product spaces by identifying their boundaries, again with a possible intermediate layering of tetrahedra.

Definition 3.3.28 (Square Product Pair) A *square product pair* is a triangulation formed using the following construction. As with square surface bundles, square product pairs come in a variety of flavours each of which is described in detail at a later point in the definition.

Let S_1 and S_2 be a pair of three-tetrahedron square twisted product spaces as described in Definition 3.2.4 with homeomorphic boundary surfaces. From Lemma 3.2.6 we must have one of the following cases.

- The boundaries of S_1 and S_2 can be tori, in which case each product space must be either a square $T^2 \times I$ or a square orientable $K^2 \times I$;
- The boundaries of S_1 and S_2 can be Klein bottles, in which case each product space must be a square non-orientable $K^2 \times I$ of type I or II.

Our aim is to identify the boundary surfaces of S_1 and S_2 according to some homeomorphism. Assume then that the boundary squares of S_1 and S_2 are arranged as illustrated in Figure 3.8 on page 95. Let α_1 , β_1 , α_2 and β_2 be the directed edges on the two boundary squares illustrated in Figure 3.34. Once more we see that in all possible cases the curves α_1 and β_1 generate the fundamental group of the boundary of S_1 and the curves α_2 and β_2 generate the fundamental group of the boundary of S_2 .




Figure 3.34: Edges α_1 , β_1 , α_2 and β_2 on the boundaries of two square product spaces

Let p , q , r and s be integers. The homeomorphism with which we identify the two boundary surfaces must identify directed edge α_1 with the curve $\alpha_2^p \beta_2^q$ and directed edge β_1 with the curve $\alpha_2^r \beta_2^s$. As with square surface bundles, this identification may require the layering of tetrahedra upon the boundary of S_2 , although such a layering is not seen in any of the triangulations presented in this chapter. Note also that there are values for p , q , r and s for which no such homeomorphism exists.

We denote the resulting triangulation by the symbol $Q_{t_1, t_2 | p, q | r, s}$, where t_1 and t_2 are symbols representing the specific square twisted product spaces S_1 and S_2 and have the following meanings.

- The symbol T represents a square $T^2 \times I$;
- The symbols K_I and K_{II} represent a square non-orientable $K^2 \times I$ of type I and type II respectively;
- The symbol K_0 represents a square orientable $K^2 \times I$.

Of the possible square product pairs, it can be observed that the only orientable triangulations are those of the form $Q_{K_0, K_0 | p, q | r, s}$.

Example 3.3.29 Consider the orientable triangulation $Q_{K_0, K_0|-1, 0|1, 1}$. This is formed from a pair of square orientable $K^2 \times I$ triangulations as illustrated in the left hand diagram of Figure 3.35. The boundary edges α_1 , α_2 , β_1 and β_2 are also illustrated in the diagram.




Figure 3.35: Constructing the square product pair $Q_{K_0, K_0|-1, 0|1, 1}$

The required homeomorphism between the two boundary tori must identify directed edges α_1 and β_1 with the curves α_2^{-1} and $\alpha_2\beta_2$ respectively. This homeomorphism can be realised by identifying face ACD with face WZY and identifying face ABD with face XWZ , as illustrated in the right hand diagram of Figure 3.35.

Theorem 3.3.30 The orientable square product pair $Q_{K_0, K_0|p, q|r, s}$ is a triangulation of the 3-manifold

$$K^2 \times I \cup K^2 \times I / \begin{bmatrix} p & r \\ q & s \end{bmatrix}.$$

Proof This result can be established simply by matching curves between the two boundary surfaces as described in Section 3.1.4. ■

The identification of the underlying 3-manifold of a non-orientable square product pair is more complex and is not presented here. Nevertheless, the underlying 3-manifold for each square product pair formed from six-tetrahedra is listed in the tables of Appendix A.

3.4 Closed Orientable Triangulations

We turn now to a presentation of all closed orientable prime minimal 3-manifold triangulations formed from at most six tetrahedra. It can be observed that each of these triangulations is of a Seifert fibred space. The triangulations are arranged into sections according to the number of tetrahedra used. Within each section triangulations are arranged according to the method of construction.

Many prime 3-manifolds allow for more than one minimal triangulation; in such cases every minimal triangulation is presented. Note however that if different minimal triangulations of a 3-manifold involve different methods of construction then these triangulations are not presented together. For a list of all closed prime minimal triangulations sorted by their underlying 3-manifolds, see the tables in Appendix A. For a brief summary of orientable census statistics, see Section 3.4.7.

Each triangulation presented here is assigned a unique name so that the triangulation can be identified within the tables in Appendices A and B. Note that the name of a triangulation is not simply the name of its underlying 3-manifold, since some 3-manifolds have many minimal triangulations. For the purpose of cross-referencing these results against those of Matveev [29], the tables in Appendix A include the index that Matveev assigns to each underlying 3-manifold.

3.4.1 One Tetrahedron

There are four single-tetrahedron orientable triangulations that appear in the census. The specific triangulations are as follows.

- The layered lens spaces $L_{1,0}$, $L_{4,1}$ and $L_{5,2}$ as described by Algorithm 3.3.7;
- The untwisted layered loop C_1 as described by Definition 3.3.8.

Note that the triangulations $L_{1,0}$ and C_1 are the one-vertex and two-vertex minimal triangulations of the 3-sphere. Note also that the triangulation $L_{4,1}$ is in fact identical to the twisted layered loop \tilde{C}_1 as described by Definition 3.3.9.

3.4.2 Two Tetrahedra

Ten different orientable triangulations are found that use two tetrahedra. These include two-vertex triangulations of $\mathbb{R}P^3$ and $L(3,1)$, which are the final two-vertex triangulations to appear in the census. The ten triangulations are as follows.

- The layered lens spaces $L_{0,1}$, $L_{2,1}$, $L_{3,1}^{(1)}$, $L_{3,1}^{(2)}$, $L_{5,1}$, $L_{7,2}$ and $L_{8,3}$ as described by Algorithm 3.3.7, where the distinction between $L_{3,1}^{(1)}$ and $L_{3,1}^{(2)}$ is explained below;
- The untwisted layered loop C_2 as described by Definition 3.3.8, which is the two-vertex triangulation of $\mathbb{R}P^3$ mentioned above;
- The twisted layered loop \tilde{C}_2 as described by Definition 3.3.9;
- The triangulation $L'_{3,1}$ illustrated in Figure 3.36. This two-vertex triangulation is formed by constructing a two-tetrahedron triangular pillow and then identifying the top and bottom faces under a 120° rotation. The underlying 3-manifold of this triangulation is the lens space $L(3,1)$.




Figure 3.36: The triangulation $L'_{3,1}$ of the lens space $L(3,1)$

It can be observed that construction described by Algorithm 3.3.7 of the layered lens space $L(3,1)$ is not unique. Specifically, the construction requires a Möbius band to be joined to an

LST(1, 1, -2) so that the boundary of the Möbius band is joined to an edge of the layered solid torus with parameter 1. There are however two such edges with parameter 1, one with degree one and one with degree five in the underlying triangulation. Joining the boundary of the Möbius band to each of these edges creates the triangulation $L_{3,1}^{(1)}$ or $L_{3,1}^{(2)}$ respectively. These two triangulations are not isomorphic.

3.4.3 Three Tetrahedra

For three tetrahedra we continue to observe only layered lens spaces and layered loops. Every triangulation from this point onwards has only one vertex, as expected from Theorem 2.4.9. The seven triangulations found are as follows.

- The layered lens spaces $L_{6,1}$, $L_{9,2}$, $L_{10,3}$, $L_{11,3}$, $L_{12,5}$ and $L_{13,5}$ as described by Algorithm 3.3.7;
- The twisted layered loop \tilde{C}_3 as described by Definition 3.3.9.

3.4.4 Four Tetrahedra

Once we reach four tetrahedra we witness the appearance of several new types of triangulation. This portion of the census includes a layered chain pair as well as two augmented triangular solid tori and a chained triangular solid torus. The following 15 triangulations are observed.

- The layered lens spaces $L_{7,1}$, $L_{11,2}$, $L_{13,3}$, $L_{14,3}$, $L_{15,4}$, $L_{16,7}$, $L_{17,5}$, $L_{18,5}$, $L_{19,7}$ and $L_{21,8}$ as described by Algorithm 3.3.7;
- The twisted layered loop \tilde{C}_4 as described by Definition 3.3.9;
- The layered chain pair $C_{2,2}$ as described by Definition 3.3.12;
- The augmented triangular solid tori $A_{2,-3}$ and $A_{3,-2}$ as described by Definition 3.3.15;
- The chained triangular solid torus $J_{1|2,-1}$ as described by Definition 3.3.19.

3.4.5 Five Tetrahedra

At the five tetrahedron level of the census we observe our first two plugged triangular solid tori. In addition the families of triangulations that have already been encountered continue to be well represented. The 40 triangulations formed from five tetrahedra are as follows.

- The layered lens spaces $L_{8,1}$, $L_{13,2}$, $L_{16,3}$, $L_{17,3}$, $L_{17,4}$, $L_{19,4}$, $L_{20,9}$, $L_{22,5}$, $L_{23,5}$, $L_{23,7}$, $L_{24,7}$, $L_{25,7}$, $L_{25,9}$, $L_{26,7}$, $L_{27,8}$, $L_{29,8}$, $L_{29,12}$, $L_{30,11}$, $L_{31,12}$ and $L_{34,13}$ as described by Algorithm 3.3.7;
- The twisted layered loop \tilde{C}_5 as described by Definition 3.3.9;
- The layered chain pair $C_{2,3}$ as described by Definition 3.3.12;
- The augmented triangular solid tori $A_{2,-5}$, $A_{3,-5}$, $A_{3,-4}$, $A_{4,-3}$, $A_{5,-3}$, $A_{5,-2}$, $A_{2,-3|2,-3}$, $A_{2,-3|3,-2}$, $A_{2,-3|3,-1}$, $A_{3,-2|3,-2}$ and $A_{3,-2|3,-1}$ as described by Definition 3.3.15;

- The chained triangular solid torus $X_{2|2,-1}$ as described by Definition 3.3.18;
- The chained triangular solid tori $J_{1|2,-3}$, $J_{1|3,-2}$, $J_{1|3,-1}$ and $J_{2|2,-1}$ as described by Definition 3.3.19;
- The plugged triangular solid tori P_0 and P'_0 as described by Definition 3.3.22.

3.4.6 Six Tetrahedra

We conclude our orientable census with the 115 six-tetrahedron triangulations listed below. Each of the families of triangulations seen thus far continues to offer representatives, and furthermore we observe our first square torus bundles and square twisted product pairs.

Six of the triangulations in the following list do not belong to any of the families of triangulations discussed in Section 3.3, although their underlying 3-manifolds can still be identified. These are the triangulations $O_{6,1}, \dots, O_{6,6}$ and are discussed in further detail below.

The 115 six-tetrahedron triangulations are then as follows.

- The layered lens spaces $L_{9,1}, L_{15,2}, L_{19,3}, L_{20,3}, L_{21,4}, L_{23,4}, L_{24,5}, L_{24,11}, L_{27,5}, L_{28,5}, L_{29,9}, L_{30,7}, L_{31,7}, L_{31,11}, L_{32,7}, L_{33,7}, L_{33,10}, L_{34,9}, L_{35,8}, L_{36,11}, L_{37,8}, L_{37,10}, L_{39,14}, L_{39,16}, L_{40,11}, L_{41,11}, L_{41,12}, L_{41,16}, L_{43,12}, L_{44,13}, L_{45,19}, L_{46,17}, L_{47,13}, L_{49,18}, L_{50,19}$ and $L_{55,21}$ as described by Algorithm 3.3.7;
- The twisted layered loop \tilde{C}_6 as described by Definition 3.3.9;
- The layered chain pair $C_{3,3}$ as described by Definition 3.3.12;
- The augmented triangular solid tori $A_{2,-7}, A_{3,-8}, A_{3,-7}, A_{4,-7}, A_{4,-5}, A_{5,-8}, A_{5,-7}, A_{5,-4}, A_{7,-5}, A_{7,-4}, A_{7,-3}, A_{7,-2}, A_{8,-5}, A_{8,-3}, A_{2,-5|2,-3}, A_{2,-3|3,-4}, A_{2,-5|3,-2}, A_{2,-5|3,-1}, A_{2,-3|3,-5}, A_{2,-3|4,-3}, A_{2,-3|4,-1}, A_{2,-3|5,-3}, A_{2,-3|5,-2}, A_{3,-5|3,-2}, A_{3,-5|3,-1}, A_{3,-4|3,-2}, A_{3,-4|3,-1}, A_{3,-2|4,-3}, A_{3,-2|4,-1}, A_{3,-2|5,-3}, A_{3,-2|5,-2}, A_{3,-1|4,-3}, A_{3,-1|5,-3}, A_{3,-1|5,-2}, A_{2,-3|2,-3|2,-3}, A_{2,-3|2,-3|3,-2}, A_{2,-3|2,-3|3,-1}, A_{2,-3|3,-2|3,-2}, A_{2,-3|3,-2|3,-1}, A_{2,-3|3,-1|3,-1}, A_{3,-2|3,-2|3,-2}, A_{3,-2|3,-2|3,-1}, A_{3,-2|3,-1|3,-1}$ and $A_{3,-1|3,-1|3,-1}$ as described by Definition 3.3.15;
- The chained triangular solid tori $X_{2|2,-3}, X_{2|3,-2}, X_{2|3,-1}$ and $X_{3|2,-1}$ as described by Definition 3.3.18;
- The chained triangular solid tori $J_{1|2,-5}, J_{1|3,-5}, J_{1|3,-4}, J_{1|4,-3}, J_{1|4,-1}, J_{1|5,-3}, J_{1|5,-2}, J_{2|2,-3}, J_{2|3,-2}, J_{2|3,-1}$ and $J_{3|2,-1}$ as described by Definition 3.3.19;
- The plugged triangular solid torus P_1 as described by Definition 3.3.22;
- The square torus bundles $T_{I|-1,0|0,-1}, T_{I|-1,1|-1,0}, T_{I|1,0|0,1}, T_{II|-1,0|1,-1}, T_{II|0,1|-1,0}$ and $T_{II|1,-1|0,1}$ as described by Definition 3.3.25;
- The square product pairs $Q_{K_0,K_0|-1,0|1,1}, Q_{K_0,K_0|0,-1|-1,0}, Q_{K_0,K_0|0,1|-1,-1}, Q_{K_0,K_0|1,0|0,1}$ and $Q_{K_0,K_0|1,1|0,-1}$ as described by Definition 3.3.28;

- The uncategorised orientable triangulations $O_{6,1}$, $O_{6,2}$, $O_{6,3}$, $O_{6,4}$, $O_{6,5}$ and $O_{6,6}$. Each of these triangulations belongs to none of the families described in Section 3.3. The underlying 3-manifolds of these triangulations are nevertheless known, since each of these triangulations can be converted through some sequence of elementary moves into a one of the categorised triangulations listed above. Elementary moves are described in more detail in Section 2.4.1 and in particular have no effect on the underlying 3-manifold.

Table 3.2 lists the underlying 3-manifold for each triangulation O_i , as well as another triangulation from the above list into which O_i can be converted through a sequence of elementary moves.

Triangulation	3-Manifold	Equivalent Triangulation
$O_{6,1}$	$S^3/P_{120} \times \mathbb{Z}_7$	$A_{3,-1 5,-2}$
$O_{6,2}$	$T^2 \times I / \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}$	$T_{1 -1,0,0,-1}$
$O_{6,3}$	$K^2 \rtimes I \cup K^2 \rtimes I / \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$	$Q_{K_0,K_0 0,-1 -1,0}$
$O_{6,4}$	$K^2 \rtimes I \cup K^2 \rtimes I / \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$	$Q_{K_0,K_0 0,-1 -1,0}$
$O_{6,5}$	$K^2 \rtimes I \cup K^2 \rtimes I / \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$	$Q_{K_0,K_0 0,-1 -1,0}$
$O_{6,6}$	$K^2 \rtimes I \cup K^2 \rtimes I / \begin{bmatrix} -1 & 1 \\ -1 & 0 \end{bmatrix}$	$Q_{K_0,K_0 0,1 -1,-1}$

Table 3.2: The uncategorised orientable triangulations $O_{6,1}, \dots, O_{6,6}$

3.4.7 Summary

We conclude the orientable census results with a brief summary of statistics. In Table 3.3 we list for each number of tetrahedra the total number of closed orientable prime minimal triangulations, as well as the number of distinct 3-manifolds that these triangulations represent. In Table 3.4 we break these total numbers of triangulations into totals for each of the families of triangulations described in Section 3.3.

Tetrahedra	Triangulations	3-Manifolds
1	4	3
2	10	7
3	7	7
4	15	14
5	40	31
6	115	74

Table 3.3: Distinct triangulations and 3-manifolds in the closed orientable census

3.5 Closed Non-Orientable Triangulations

The census of closed non-orientable prime minimal 3-manifold triangulations is simpler than the corresponding orientable census since there are in fact only three such non-orientable triangulations

Tetrahedra	1	2	3	4	5	6
Layered lens spaces	3	7	6	10	20	36
Untwisted layered loops	1	1	0	0	0	0
Twisted layered loops	0	1	1	1	1	1
Layered chain pairs	0	0	0	1	1	1
Augmented tri. solid tori	0	0	0	2	11	44
Chained tri. solid tori (axial)	0	0	0	0	1	4
Chained tri. solid tori (major)	0	0	0	1	4	11
Plugged tri. solid tori (major)	0	0	0	0	1	1
Plugged tri. solid tori (minor)	0	0	0	0	1	0
Square surface bundles	0	0	0	0	0	6
Square product pairs	0	0	0	0	0	5
Miscellaneous	0	1	0	0	0	6
Total	4	10	7	15	40	115

Table 3.4: Frequencies of orientable triangulations from different families

formed from at most five tetrahedra. A significant number of six-tetrahedron triangulations have been found however; each of these represents either a torus bundle or a Klein bottle bundle over the circle.

As with the orientable census results, the triangulations are arranged by number of tetrahedra and then by method of construction. Most non-orientable 3-manifolds that appear in this census have several minimal triangulations; for each such 3-manifold every minimal triangulation is presented. For lists of triangulations that are arranged by underlying 3-manifold instead of by method of construction, the reader is again referred to the tables in Appendix A. A summary of non-orientable census statistics is offered in Section 3.5.3.

As with the orientable triangulations, each non-orientable triangulation presented here is assigned a unique name for the purposes of identifying the triangulation in Appendices A and B.

3.5.1 Five or Fewer Tetrahedra

There are in fact just three closed non-orientable prime minimal triangulations containing at most five tetrahedra. Specifically there is a single two-tetrahedron triangulation of the non-orientable 2-sphere bundle over the circle $S^2 \tilde{\times} S^1$, and there are two three-tetrahedron triangulations of the product $\mathbb{R}P^2 \times S^1$.

- The two-tetrahedron triangulation of the non-orientable 2-sphere bundle $S^2 \tilde{\times} S^1$ is illustrated in Figure 3.37. This triangulation is formed from two tetrahedra joined along common faces PRQ and PSQ . To complete the triangulation we identify the upper face PRS with the lower face SQR and we identify the lower face PRS with the upper face SQR .

We denote this specific triangulation by the symbol N_2 .

- The two three-tetrahedron triangulations of the product $\mathbb{R}P^2 \times S^1$ are illustrated in Figure 3.38. In the first diagram, faces BFE and AED are identified and faces BCE and ABD are identified, forming a solid torus (and in fact a layered chain of length 3). Upper faces




Figure 3.37: The minimal triangulation of the non-orientable 2-sphere bundle $S^2 \tilde{\times} S^1$

AFB and EAF are then identified with lower faces DCE and BDC respectively, converting this orientable 2-ball bundle $B^2 \times S^1$ into the non-orientable product $\mathbb{R}P^2 \times S^1$.



Figure 3.38: The minimal triangulations of $\mathbb{R}P^2 \times S^1$

In the second diagram, faces BFE and DEA are identified and faces BCE and DBA are identified, forming a solid Klein bottle. Upper faces AFB and EAF are then identified with lower faces DCE and BDC respectively, converting this non-orientable 2-ball bundle $B^2 \tilde{\times} S^1$ into the product $\mathbb{R}P^2 \times S^1$.

We denote the first and second of these triangulations by the symbols $N_{3,1}$ and $N_{3,2}$ respectively. Note that it is possible to construct $N_{3,2}$ by identifying the two faces of a square $T^2 \tilde{\times} I$ as described by Definition 3.2.5.

3.5.2 Six Tetrahedra

Of the 24 closed non-orientable prime minimal triangulations formed from six tetrahedra, 20 can be classified as either square surface bundles or square product pairs. The remaining four triangulations do not belong to any of the families described in Section 3.3, and are discussed in further detail below.

Although we present 24 distinct triangulations, these in fact only represent five distinct 3-manifolds, all of which are either torus bundles or Klein bottle bundles over the circle. The individual triangulations are as follows.

- The square torus bundles $T_{I|-1,0|-1,1}$, $T_{I|0,-1|-1,0}$, $T_{I|0,1|1,0}$, $T_{I|1,0|1,-1}$, $T_{II|-1,1|1,0}$ and $T_{II|1,0|0,-1}$ as described by Definition 3.3.25;
- The square Klein bottle bundles $K_{I|-1,0|-1,-1}$, $K_{I|-1,0|-1,1}$, $K_{I|1,0|0,-1}$, $K_{I|1,0|0,1}$, $K_{II|-1,0|0,-1}$, $K_{II|0,-1|-1,0}$, $K_{II|0,1|1,0}$, $K_{II|1,0|0,1}$, $K_{III|-1,0|0,-1}$, $K_{III|0,-1|-1,0}$, $K_{III|0,1|1,0}$ and $K_{III|1,0|0,1}$ as described by Definition 3.3.25, where three of these square Klein bottle bundles are isomorphic to square torus bundles as discussed below;
- The square product pairs $Q_{T,T|0,1|-1,-1}$, $Q_{T,K_0|-1,-1|1,0}$, $Q_{T,K_0|0,1|-1,-1}$, $Q_{T,K_0|1,0|0,1}$ and $Q_{K_{II},K_{II}|1,0|0,1}$ as described by Definition 3.3.28;

- The uncategorised non-orientable triangulations $N_{6,1}$, $N_{6,2}$, $N_{6,3}$ and $N_{6,4}$. Each of these triangulations belongs to none of the families listed in Section 3.3, but can nevertheless be converted through a sequence of elementary moves into the triangulation $K_{\text{II}|0,1|1,0}$. Using Theorem 3.3.27 it can thus be seen that the underlying 3-manifold of each of these triangulations is the Klein bottle bundle $K^2 \times I / \begin{bmatrix} -1 & 1 \\ 0 & -1 \end{bmatrix}$.

Although 27 symbolic names for triangulations are presented in the list above, three of the square torus bundles and three of the square Klein bottle bundles are isomorphic in pairs. Specifically, triangulations $T_{\text{I}|-1,0|-1,1}$ and $K_{\text{I}|1,0|0,1}$ are isomorphic, $T_{\text{I}|0,1|1,0}$ and $K_{\text{II}|0,-1|-1,0}$ are isomorphic and $T_{\text{II}|1,0|0,-1}$ and $K_{\text{I}|-1,0|-1,1}$ are isomorphic, leaving 24 unique triangulations in our list.

3.5.3 Summary

Once more we closed our census results with a short statistical summary. Table 3.5 lists for each number of tetrahedra the total number of closed non-orientable prime minimal triangulations, as well as the number of distinct 3-manifolds that these triangulations represent.

Tetrahedra	Triangulations	3-Manifolds
1	0	0
2	1	1
3	2	1
4	0	0
5	0	0
6	24	5

Table 3.5: Distinct triangulations and 3-manifolds in the closed non-orientable census

3.6 Normal Surfaces

As described in Section 1.1, the normal surfaces within a 3-manifold triangulation offer insight into the properties of the triangulation and of the underlying 3-manifold. The vertex normal surfaces within a triangulation are however notoriously difficult to calculate by hand, and very little software is available to automate this task.

Therefore tables are provided in Appendix B that describe the vertex normal surfaces of every closed prime minimal triangulation formed from at most six tetrahedra. This enumeration is performed under both the standard triangle-quadrilateral coordinate system and the quadrilateral-only coordinate system of Tollefson [45]. Once more all calculations were performed using the program *Regina*.

Although space constraints do not allow for a presentation of the full vector corresponding to each normal surface, a variety of properties of each surface is presented. These properties include Euler characteristic, orientability, whether the surface is one-sided or two-sided, whether the surface is an edge link and whether the surface is a splitting surface. Splitting surfaces are discussed in detail in Chapter 4.

Note that for the purposes of identifying edge links, only thin edge links are recognised, i.e., surfaces that can be represented as the boundary of a small neighbourhood of an edge. Thick edge links as described by Jaco and Rubinstein [18] are not explicitly marked as such.

3.7 Future Directions

As discussed in the opening comments of this chapter, of the 218 distinct orientable and non-orientable triangulations listed in sections 3.4 and 3.5, all but 14 can be described as members of the families presented in Section 3.3. Of the remaining 14 triangulations, four are the trivial triangulations $L'_{3,1}$, N_2 , $N_{3,1}$ and $N_{3,2}$. The remaining ten triangulations $O_{6,1}, \dots, O_{6,6}$ and $N_{6,1}, \dots, N_{6,4}$ have yet to be categorised into families sharing a common large-scale structure.

It is difficult to generalise the combinatorial structures of such a small number of remaining triangulations. It is hoped that as the results of the seven-tetrahedron census are analysed, additional families of triangulations can be found that encompass the uncategorised six-tetrahedron triangulations.

This same process of using the results of an $(n+1)$ -tetrahedron census to complete the categorisation of an n -tetrahedron census has been used repeatedly in developing the results presented in this chapter. For instance, the layered chain pair construction was only generalised when processing the five-tetrahedron census even though the layered chain pair $C_{2,2}$ appears in the four-tetrahedron census. Similarly, the generalised form of the orientable square torus bundles was not understood until the non-orientable census results were processed.

The computationally intensive portions of the seven-tetrahedron and eight-tetrahedron orientable census generation are already complete at the time of writing. The subsequent processing of triangulations is still underway, and it is hoped that further general families of triangulations can be described as a result.

Additional research is progressing into the analysis of normal surfaces within these general families of triangulations. It can be proven for instance that, when calculating normal surfaces using quadrilateral-only coordinates, the twisted layered loop \tilde{C}_n has precisely $n + 1$ vertex normal surfaces, each of which corresponds to the link of one of the $n + 1$ edges of the triangulation. The link of the hinge edge of \tilde{C}_n is furthermore the double of a one-sided Klein bottle which forms a splitting surface within the triangulation. These results can indeed be observed in the tables of Appendix B.

In a similar fashion it is hoped to prove general results for other families of triangulations regarding the normal surfaces within these triangulations. The tables of surfaces listed in Appendix B provide a source of inspiration in this endeavour.

Chapter 4

Splitting Surfaces

Recall from Section 1.1 that the recognition and analysis of normal surfaces within a 3-manifold can lead to a greater understanding of this 3-manifold. In this chapter we study splitting surfaces, a particular class of normal surfaces that contain only quadrilateral discs.

Splitting surfaces have a number of properties that justify our interest in them. They realise a Heegaard splitting of the underlying 3-manifold and they are algorithmically easy to find. In addition, the quadrilateral structure of a splitting surface can be used to uniquely reconstruct the 3-manifold triangulation in which it is contained. Each of these results is presented in Section 4.1. Furthermore, 2-dimensional surfaces and their cell decompositions are significantly more well understood than 3-manifolds and their triangulations. For all of these reasons splitting surfaces offer themselves as a useful tool in the analysis of triangulations of 3-manifolds.

In section 4.2 we examine in more detail the quadrilateral structures of splitting surfaces. Section 4.3 presents splitting surface signatures, a computation tool used to assist with the analysis and generation of splitting surfaces. The analysis of signatures is then discussed in Section 4.4 and in Section 4.5 we examine a census of all signatures of order ≤ 8 and their corresponding triangulations.

Research into splitting surfaces is still in progress. While this chapter develops a basic infrastructure for studying splitting surfaces, it is hoped that with further research the tools for analysing these surfaces can be further enriched. Possible directions for such future research are discussed in Section 4.6.

Note that in the lists of normal surfaces presented in Appendix B for the orientable and non-orientable censuses of 3-manifold triangulations, splitting surfaces are identified where they appear.

4.1 Motivation

We open this section with a discussion of the basic ideas behind splitting surfaces. Following this we prove a series of results that, as discussed in the opening comments of this chapter, explain why splitting surfaces in particular are worthy of investigation.

Definitions 4.1.1 (Splitting Surfaces) A *splitting surface* within a 3-manifold triangulation T is a normal surface consisting of precisely one quadrilateral disc within each tetrahedron of T and

no other normal discs.

The *hinge edges* of a triangulation with respect to a splitting surface are the edges of the triangulation that the splitting surface does not intersect.

Example 4.1.2 Consider the two-tetrahedron triangulation of the 3-manifold S^3/Q_8 as illustrated in the left hand diagram of Figure 4.1. This triangulation is closed; the three faces at the front of the diagram are identified in various ways with the three faces at the rear of the diagram. The different arrowheads illustrate the corresponding edge identifications.




Figure 4.1: A splitting surface within S^3/Q_8

The two quadrilaterals shaded in the right hand diagram of this figure form a splitting surface within this triangulation. The resulting identifications of the quadrilateral edges are again illustrated by the different arrowheads. It can be seen that these two quadrilaterals together form a Klein bottle. Edge e is the only hinge edge corresponding to this splitting surface.

Lemma 4.1.3 Any splitting surface within a single 3-manifold triangulation is connected.

Proof If two tetrahedra meet along a common face in the triangulation then the quadrilateral discs within these tetrahedra must meet along a common edge, since each tetrahedron contains only one normal disc and this disc meets all four faces of the tetrahedron. It follows then that since the triangulation is connected the splitting surface formed from these normal discs must also be connected. ■

Theorem 4.1.4 A one-sided or two-sided splitting surface within a closed triangulation represents a one-sided or two-sided Heegaard splitting respectively of the underlying 3-manifold.

Furthermore, consider the set of all hinge edges in the triangulation; these hinge edges together form one or more graphs within the 1-skeleton of the triangulation. Each handlebody into which the 3-manifold is decomposed by this Heegaard splitting has precisely one of these graphs at its core, and conversely each such graph is the core of precisely one of these handlebodies.

Proof If we cut the 3-manifold along the splitting surface, each tetrahedron is split into two triangular prisms as illustrated in the left hand diagram of Figure 4.2. Each prism meets the resulting boundary surface along precisely one quadrilateral face and nowhere else; these quadrilateral faces are shaded in the diagram. Note also that each triangular prism meets precisely one hinge edge.

The neighbourhood of each hinge edge is thus a ring of prisms as illustrated in the right hand diagram of Figure 4.2. These rings of prisms meet each other along the triangular faces at their ends, and each such meeting induces a meeting of the hinge edges at their cores.




Figure 4.2: Splitting tetrahedra into triangular prisms

It follows that each connected component of the resulting structure is a handlebody with a hinge edge graph at its core, and that each such graph forms the core of precisely one of these handlebodies. Furthermore since each handlebody has the splitting surface as its boundary and this splitting surface is connected by Lemma 4.1.3, we see that the splitting surface represents a Heegaard splitting as claimed. ■

Theorem 4.1.5 A splitting surface within a 3-manifold triangulation is both a fundamental normal surface and a vertex normal surface as described in Section 1.1.

Proof Let S be a splitting surface within some triangulation. If S is not a vertex normal surface then there are normal surfaces S_1 and S_2 neither of which is a multiple of S and for which $kS = lS_1 + mS_2$ for some positive integers k, l and m .

Consider the surface S_1 . Since S contains no triangular discs it follows that S_1 likewise contains no triangular discs. Therefore every disc of S_1 is a quadrilateral and thus meets all four faces of its enclosing tetrahedron. It follows from the normal surface matching equations that if Δ_1 and Δ_2 are adjacent tetrahedra in our triangulation then S_1 must have the same number of quadrilateral discs in each of Δ_1 and Δ_2 . Since our triangulation is connected we can conclude that S_1 contains precisely q quadrilateral discs in every tetrahedron for some constant q .

Finally since $kS = lS_1 + mS_2$ we see that the quadrilaterals of S_1 within each tetrahedron must be of the same type as the quadrilaterals of S . Therefore $S_1 = qS$, contradicting our earlier claim that S_1 is not a multiple of S . It follows that S is a vertex normal surface, and since S is not an integer multiple of some other surface we see that S is also a fundamental normal surface. ■

Theorem 4.1.5 is particularly appealing because it provides us with an easy method of enumerating all of the splitting surfaces within a triangulation. We construct the vertex normal surfaces as described in Section 1.1, and then simply examine each one to see whether it contains precisely one quadrilateral for each tetrahedron and no other normal discs.

Theorem 4.1.6 Let Q be a quadrilateral decomposition of some surface. Then there is at most one 3-manifold triangulation (up to isomorphism) that contains Q as a splitting surface.

Proof We prove this by constructing the enclosing 3-manifold triangulation. We begin by taking each individual quadrilateral of Q and enclosing it within a tetrahedron. The faces of these tetrahedra are then identified as follows.

Let f be a face of some tetrahedron Δ and let q be the quadrilateral disc within Δ . Then precisely one edge e of q runs along face f . If e is a boundary edge of the surface Q then we declare

f to be a boundary face of the overall triangulation. Otherwise edge e is identified with some other edge e' in our quadrilateral decomposition. Let e' belong to quadrilateral q' within tetrahedron Δ' , and let e' run along face f' of tetrahedron Δ' . Then we identify face f of Δ with face f' of Δ' in the unique way that allows edges e and e' to be identified correctly. This procedure is illustrated in Figure 4.3.




Figure 4.3: Reconstructing a triangulation from a splitting surface

Note that the face identifications described above might result in a structure that is not actually a 3-manifold triangulation. Some tetrahedron vertices might have links that are neither 3-spheres nor discs, and some tetrahedron edges might be identified with themselves in reverse.

Nevertheless, each step in the above construction is necessary for Q to be a splitting surface within the resulting triangulation. Furthermore this construction determines the identification (or lack of identification) for every face of every tetrahedron. Thus if a 3-manifold triangulation does exist that contains Q as a splitting surface then this 3-manifold triangulation is unique. ■

Note that not only does Theorem 4.1.6 show that a 3-manifold can be reconstructed from one of its splitting surfaces, but the proof also presents a straightforward algorithm for performing this reconstruction.

As an aside, if a 3-manifold triangulation T does contain a splitting surface S then the dual 1-skeleton of S is in fact the face pairing graph of T as defined in Section 2.6.1. This can be seen from the fact that adjacent quadrilaterals of S correspond precisely to adjacent tetrahedra of T .

4.2 Diamond Representations

As observed in the proof of Theorem 4.1.6, not every quadrilateral decomposition of a surface corresponds to a splitting surface within some 3-manifold triangulation. We aim here to impose additional structure upon such quadrilateral decompositions so that potential splitting surfaces are more easily recognised.

Definition 4.2.1 (Diamond Representation) A *diamond representation* of a splitting surface within a 3-manifold is constructed as follows. We begin by assigning an arbitrary orientation to each hinge edge. Observing that each tetrahedron has precisely two edges that are hinge edges, we arbitrarily declare one of these edges to be the *front hinge* and the other to be the *rear hinge* for each tetrahedron. Note that the front and rear hinges might or might not be identified in the overall triangulation.

Each tetrahedron can now be drawn as a square with the front hinge running above the square from bottom left to top right and the back hinge running behind the square from bottom right to top left. This is illustrated for a single tetrahedron in Figure 4.4. The quadrilateral within this

tetrahedron appears as a diamond within the square; we refer to this presentation as a diamond representation of the quadrilateral. Note that once the hinge edges have been oriented and the front and rear hinges have been selected for each tetrahedron, the diamond representation for each quadrilateral is uniquely determined.




Figure 4.4: A diamond representation of a single quadrilateral

The diamond representations of every quadrilateral in the splitting surface together form a diamond representation of the splitting surface. Note that a splitting surface may have many different diamond representations according to the different ways in which the above choices can be made.

Example 4.2.2 Consider the splitting surface of S^3/Q_8 illustrated in Figure 4.1 on page 128. The single hinge edge e is oriented as illustrated in the diagram. For each tetrahedron we declare the instance of e at the top of the tetrahedron to be the front hinge and the instance of e at the base of the tetrahedron to be the rear edge. The resulting diamond representation of the splitting surface is illustrated in Figure 4.5.




Figure 4.5: A diamond representation of a splitting surface within S^3/Q_8

Lemma 4.2.3 Consider some diamond representation of a splitting surface and let q_1 and q_2 be quadrilaterals in this diamond representation that are adjacent along some common edge e . Note that q_1 and q_2 may in fact be the same quadrilateral.

There are four possible locations for edge e in quadrilateral q_1 , four possible locations for edge e in quadrilateral q_2 and two possible orientations with which these edges may be identified, leading to a total of 32 possible ways in which q_1 and q_2 may be identified along a common edge.

Of these 32 possible edge identifications, 16 are represented in Figure 4.6. These 16 edge identifications can be seen by selecting any of the four edges of q_1 on the left side of the diagram and identifying it with the edge marked with the matching arrowhead in any of the four possible quadrilaterals q_2 on the right side of the diagram. The directions of the arrowheads indicate the orientation associated with each identification.

It is then true that in a diamond representation of a splitting surface within a 3-manifold triangulation, every edge identification must adhere to one of these 16 methods of identification.




Figure 4.6: The 16 allowable quadrilateral edge identifications

Proof Let quadrilaterals q_1 and q_2 belong to tetrahedra Δ_1 and Δ_2 respectively. For each tetrahedron Δ_i , let h_i denote the unique hinge edge that is parallel to edge e and that borders the tetrahedron face containing e . Since the two tetrahedron faces containing edge e are identified in the underlying triangulation, it follows as a consequence of this face identification that hinge edges h_1 and h_2 must likewise be identified.

Since both hinge edges already have an orientation assigned to them, it is necessary to ensure that h_1 and h_2 are identified in such a way that their orientations match, since these form a single hinge edge in the underlying triangulation which is assigned a single orientation. Of the 32 possible edge identifications, it can be seen that the 16 identifications illustrated in Figure 4.6 are precisely the identifications for which the hinge orientations are consistent. ■

Definitions 4.2.4 Consider the identification of two quadrilateral edges in a diamond representation of a splitting surface within a 3-manifold triangulation. If the corresponding identification of tetrahedron faces causes the front hinges of each tetrahedron to be identified and the back hinges of each tetrahedron to be identified, this edge identification is said to be *hinge-preserving*. Otherwise the corresponding identification of tetrahedron faces causes the front hinge of each tetrahedron to be identified with the back hinge of the other, and this edge identification is said to be *hinge-reversing*.

Figure 4.4 shows each tetrahedron embedded in \mathbb{R}^3 ; this embedding can be used to induce an orientation on each tetrahedron. If the corresponding identification of tetrahedron faces is an orientation-preserving map in the underlying triangulation, this edge identification is said to be *oriented*. Otherwise the corresponding identification of tetrahedron faces is an orientation-reversing map in the underlying triangulation and this edge identification is said to be *non-oriented*.

Example 4.2.5 The 16 possible edge identifications in Figure 4.6 are marked in the figure as hinge-preserving or hinge-reversing and oriented or non-oriented according to which of the four possible diagrams is used for quadrilateral q_2 . There are precisely eight hinge-preserving and eight hinge-reversing identifications, and similarly there are precisely eight oriented and eight non-oriented edge identifications.

We close this section by defining a particular class of diamond representations that are analysed in detail in the later parts of this chapter.

Definition 4.2.6 (Natural Representation) Let S be a splitting surface within some orientable 3-manifold triangulation. Assume that the hinge edges in the triangulation have been oriented and the front and rear hinges selected for each tetrahedron so that every quadrilateral edge identification in S is oriented as described in Definitions 4.2.4. Then the hinge orientations, the front and rear hinge selections and the corresponding diamond representation of S are together called a *natural representation* of S .

Example 4.2.7 The diamond representation of the splitting surface within S^3/Q_8 that is illustrated in Figure 4.5 forms a natural representation of the surface.

Note that not every splitting surface within an orientable 3-manifold triangulation allows for a natural representation as seen in the following example.

Example 4.2.8 The three-quadrilateral diamond representation of a surface illustrated in Figure 4.7 represents the only splitting surface within the minimal triangulation of the lens space $L(10, 3)$. This splitting surface does not allow for a natural representation.




Figure 4.7: A splitting surface with no natural representation

4.3 Signatures

Splitting surfaces offer a great deal of potential for the computational analysis of 3-manifolds. They are easily identified within a triangulation as seen in Theorem 4.1.5, and working with 2-dimensional surfaces is significantly easier than working with 3-manifold triangulations. To aid the computational analysis and manipulation of splitting surfaces, we introduce the concept of splitting surface signatures, a particularly compact method of representing splitting surfaces within closed orientable 3-manifold triangulations.

Although diamond representations are a step towards the standardisation of splitting surface presentations, there are still many different diamond representations for any given splitting surface. In this section we therefore use natural representations as described in Definition 4.2.6 to further standardise such presentations.

Lemma 4.3.1 In a natural representation of a splitting surface within a closed orientable 3-manifold triangulation, every edge on the upper half of a quadrilateral diamond is identified with an edge on the lower half of a quadrilateral diamond (and vice versa).

Proof Since the triangulation is closed, every quadrilateral edge must be identified with some other quadrilateral edge. The result can then be seen by simply checking all eight possible oriented




Figure 4.8: The eight oriented quadrilateral edge identifications

edge identifications as illustrated in Figure 4.6. For convenience these oriented edge identifications are reproduced in Figure 4.8. ■

We now wish to examine paths along our splitting surface that run from quadrilateral to quadrilateral. The building blocks of these paths are left arrows and right arrows as described below.

Definitions 4.3.2 (Arrows) Consider the diamond representation of a single quadrilateral q . Upon this quadrilateral draw two directed lines joining the midpoints of opposite edges as illustrated in Figure 4.9. These directed lines are called *arrows*. To distinguish between these two arrows we declare the *left arrow* to run from the bottom right edge to the top left edge of the quadrilateral, and the *right arrow* to run from the bottom left edge to the top right edge of the quadrilateral. We denote the left arrow by q^- and the right arrow by q^+ .




Figure 4.9: Left and right arrows on a quadrilateral

From Lemma 4.3.1 we see that in a natural representation of a splitting surface within a closed orientable 3-manifold triangulation, the tip of each arrow meets the base of an arrow on an adjacent quadrilateral. Thus the arrows naturally arrange themselves into cycles.

Definitions 4.3.3 (Paths) A *path* in a natural representation of a splitting surface is a cycle of arrows formed by identifying the tip of each arrow to the base of the following arrow according to the corresponding quadrilateral edge identifications.

A path that follows arrows $\alpha_1, \alpha_2, \dots, \alpha_k$ in turn is denoted $(\alpha_1 \alpha_2 \dots \alpha_k)$. Note that the tip of arrow α_k must be identified with the base of arrow α_1 .

The *length* of a path is the number of arrows that it contains. Note that a path of length k can be expressed in k different ways by selecting any of the k possible starting points.

Example 4.3.4 Consider the natural representation illustrated in Figure 4.10, which represents the same splitting surface within S^3/Q_8 that was seen in Figure 4.5. Label the quadrilaterals a and b as shown. Then there is a single path in this surface which can be expressed as $(a^+ b^+ a^- b^-)$.




Figure 4.10: Finding a path within a splitting surface of S^3/Q_8

Lemma 4.3.5 Consider the set of all paths in a natural representation of a splitting surface within a closed orientable 3-manifold triangulation formed from n tetrahedra. Each of the $2n$ possible arrows appears in precisely one of these paths.

Proof Using Lemma 4.3.1 and the fact that the underlying triangulation is closed, we see that each arrow tip meets precisely one arrow base. The result then follows. ■

Definitions 4.3.6 (Signatures) A *signature* of a natural representation of a splitting surface within a closed orientable 3-manifold triangulation is obtained as follows. We begin by writing a list of all paths within the natural representation. Following this we replace each arrow identifier q^- or q^+ with the corresponding quadrilateral identifier q . A *cycle* within this signature is a portion of the signature corresponding to a single path.

Example 4.3.7 Consider the natural representation illustrated in Figure 4.11. This natural representation has a signature $(ab)(bc)(a)(c)$. This signature contains four individual cycles corresponding to the four different paths of lengths 2, 2, 1 and 1 within this natural representation.




Figure 4.11: A splitting surface with signature $(ab)(bc)(a)(c)$

Lemma 4.3.8 In any signature, each quadrilateral identifier appears precisely twice.

Proof This follows immediately from Lemma 4.3.5 since there are precisely two arrows corresponding to each quadrilateral. ■

Theorem 4.3.9 Let σ_1 be a signature of splitting surface S_1 within the closed orientable 3-manifold triangulation T_1 and let σ_2 be a signature of splitting surface S_2 within the closed orientable 3-manifold triangulation T_2 . If σ_1 and σ_2 are identical signatures then splitting surfaces S_1 and S_2 have isomorphic quadrilateral structures and triangulations T_1 and T_2 are likewise isomorphic.

Proof Note first that the list of all paths within a natural representation provides enough information to reconstruct the natural representation completely, since the connections between arrows correspond precisely to the quadrilateral edge identifications in the corresponding diamond representation. For instance, the presence of the adjacent symbols $\dots a^+b^- \dots$ within a path indicates that the top right edge of quadrilateral a is identified with the bottom right edge of quadrilateral b , where the orientation of this edge identification is determined by Lemma 4.2.3.

Since signatures σ_1 and σ_2 are the same, we see that the sets of paths in the natural representations of splitting surfaces S_1 and S_2 differ only in that the left and right arrows q^- and q^+ may be interchanged for a variety of quadrilaterals q . Such an interchange can be produced by reflecting quadrilateral q in a vertical axis, which in turn corresponds to spinning the tetrahedron containing q by 180° about this same vertical axis.

The effect of spinning a tetrahedron by 180° about a vertical axis is that the front and rear hinges are interchanged, although their orientations are preserved and any oriented identifications of quadrilateral edges remain oriented. In particular, transforming a natural representation of a splitting surface in this way produces another natural representation of the same splitting surface.

It follows then that the natural representations of splitting surfaces S_1 and S_2 differ only by a sequence of such 180° rotations of tetrahedra, and thus S_1 and S_2 have the same underlying quadrilateral structure and are therefore isomorphic. Using Theorem 4.1.6 we then conclude that triangulations T_1 and T_2 are likewise isomorphic. ■

Thus we see that a splitting surface signature, like the quadrilateral structure of the splitting surface itself, provides enough information for the reconstruction of the underlying 3-manifold triangulation.

4.4 Analysing Signatures

Although we see from Theorem 4.3.9 that a signature uniquely defines its underlying splitting surface and 3-manifold triangulation, it is not clear from simply looking at an arbitrary signature which 3-manifold it represents. Furthermore, given an arbitrary collection of cycles formed from quadrilateral labels, it is not clear whether this collection of cycles even forms a signature of a splitting surface within some 3-manifold triangulation. In this section we therefore examine a variety of methods with which we can extract information from either a signature or a more general collection of cycles.

Definition 4.4.1 (Signatures Revised) For the remainder of this chapter we redefine a *signature* as follows. Let \mathcal{L} be some set of labels. A collection of cycles formed from elements of \mathcal{L} is called a *signature* if each element of \mathcal{L} appears precisely twice within this collection of cycles.

The *order* of the signature is the number of different labels that are used, i.e., the size $|\mathcal{L}|$. Note that the sum of the individual cycle lengths is always twice the order of the signature.

It can be observed from Lemma 4.3.8 that this revised definition is an extension of Definitions 4.3.6. That is, if σ is a signature of a natural representation of a splitting surface within a closed orientable 3-manifold triangulation as defined earlier, then σ remains a signature under this

new definition. There are however signatures under this new definition that do not correspond to a splitting surface within a 3-manifold triangulation.

Example 4.4.2 Using the set $\mathcal{L} = \{a, b, c\}$ we see that $(acc)(ab)(b)$ is a signature of order three.

Definitions 4.4.3 Let σ be a signature of order n and let S be a collection of n quadrilaterals labelled using the n labels that appear in σ . By arbitrarily replacing the two occurrences within σ of each label l with the two symbols l^- and l^+ , we can convert σ into a collection of paths formed from left arrows and right arrows on our n quadrilaterals. These paths can then be used to identify the edges of our quadrilaterals in pairs as described in the proof of Theorem 4.3.9; in particular only the 16 allowable edge identifications described in Lemma 4.2.3 are used. The result of this procedure is a quadrilateral decomposition of a closed surface. This surface and its quadrilateral decomposition are called the *underlying surface* of σ .

Let T be a collection of n tetrahedra. From the underlying surface of σ we can use the construction described in the proof of Theorem 4.1.6 to identify the faces of these tetrahedra in pairs. The resulting structure is called the *enclosing triangulation* of σ . As noted in the proof of Theorem 4.1.6, this might or might not in fact be a 3-manifold triangulation. In particular some vertices of the triangulation might have links that are neither spheres nor discs. We must also consider the possibility that some edges of the triangulation might be identified with themselves in reverse, though it will be seen in Lemma 4.4.5 that this is never the case.

Lemma 4.4.4 The underlying surface and enclosing triangulation of a signature are well-defined up to isomorphism.

Proof Let σ be some signature. By following an argument analogous to the proof of Theorem 4.3.9, we see that any two underlying surfaces of σ differ only by a sequence of reflections of quadrilaterals in vertical axes and thus have isomorphic quadrilateral decompositions. Since the construction described in Theorem 4.1.6 requires no choices to be made, it follows then that any two enclosing triangulations of σ are similarly isomorphic. ■

Lemma 4.4.5 Let T be the enclosing triangulation of a signature. Then T is orientable with no boundary faces, and no edge of T is identified with itself in reverse.

Proof Since each quadrilateral edge identification in the underlying surface corresponds to the meeting of the tip of an arrow with the base of an arrow, it follows that each such edge identification involves one of the two upper edges of a quadrilateral and one of the two lower edges of a quadrilateral. From Figure 4.6 on page 132 it can be seen that each such edge identification is oriented. Thus from Definitions 4.2.4 we see that triangulation T is orientable.

The observation that T has no boundary faces is immediate from the fact that the construction described in Definitions 4.4.3 leaves no edge of the underlying surface unmatched and therefore no face of the enclosing triangulation unmatched.

Furthermore, recall from the proof of Lemma 4.2.3 that each edge identification described in Figure 4.6 identifies the corresponding hinge edges with consistent orientations. Therefore no hinge edge in T is identified with itself in reverse. Finally, since each remaining edge of T passes through

a vertex of the underlying surface and since a small neighbourhood of any vertex in the underlying surface is two-sided, it follows that no other edge of T can be identified with itself in reverse. ■

Lemma 4.4.6 Let σ be a signature with underlying surface S and enclosing triangulation T . If we extend the concept of a normal surface to include surfaces within triangulations whose vertex links might be neither spheres nor discs, we see that S is a splitting surface within T and furthermore that σ is a signature of a natural representation of S as described by Definitions 4.3.6.

Conversely, if σ is a signature of a natural representation of the splitting surface S within the closed orientable 3-manifold triangulation T as described by Definitions 4.3.6, then the underlying surface of σ is isomorphic to S and the enclosing triangulation of σ is isomorphic to T .

Proof Let σ be a signature with underlying surface S and enclosing triangulation T as described above. It is clear from the construction of S that σ is a signature of a natural representation of S as described by Definitions 4.3.6. Similarly it is clear from the construction of T that S forms a splitting surface within T .

Conversely, let σ be a signature of a natural representation of the splitting surface S within the closed orientable 3-manifold T as described by Definitions 4.3.6. By following an identical argument to the proof of Theorem 4.3.9, we can prove that the underlying surface of σ has a quadrilateral decomposition isomorphic to that of S . We can continue with an argument identical to the proof of Theorem 4.1.6 to show that the enclosing triangulation of σ is isomorphic to T . ■

Having verified that Definitions 4.4.1 and 4.4.3 are consistent with the concepts developed earlier in this chapter, we turn to extracting properties of the enclosing triangulation from a signature.

Definition 4.4.7 (Equivalence) Two signatures are *equivalent* if their underlying surfaces have isomorphic quadrilateral decompositions.

Lemma 4.4.8 Applying any of the following operations to a signature results in an equivalent signature.

- Relabelling the quadrilaterals. For instance, $(abcabc)$ can be converted to $(bacbac)$ by switching labels a as b .
- Moving the front portion of a cycle to the rear of a cycle. For instance, $(abcabc)(dede)$ can be converted to $(cabcab)(dede)$ by moving the first two letters of the first cycle to the end of the cycle.
- Reversing all of the cycles. For instance, $(abcabc)(dede)$ can be converted to $(cbacba)(eded)$.

Proof It is clear that relabelling the quadrilaterals produces an equivalent signature. Furthermore, moving the front portion of a cycle to the rear simply corresponds to beginning at a different arrow when transcribing the corresponding path in the underlying surface.

Reversing all of the cycles within a signature can be achieved by reversing the orientation of every hinge edge in the enclosing triangulation. This in turn corresponds to rotating every quadrilateral in the underlying surface by 180° about its centre. The underlying surfaces may have different diamond representations as result, but their quadrilateral structures will remain isomorphic. ■

Note that reversing some but not all of the cycles within a signature might not produce an equivalent signature. For instance, signatures $(abcd)(bcd)$ and $(abcd)(dcb)$ can be seen to have non-isomorphic underlying surfaces.

Lemma 4.4.9 The enclosing triangulation of a signature is connected if and only if the cycles that form the signature can be partitioned into two groups so that no label appears in both groups.

Proof Let Δ and Δ' be two tetrahedra in the enclosing triangulation whose corresponding quadrilaterals have labels q and q' . It can be seen from the construction of the enclosing triangulation that Δ and Δ' are adjacent tetrahedra if and only if q and q' are adjacent symbols within some cycle of the signature.

If the cycles can be partitioned as described, this partition of cycles induces a partition of the labels and thus a partition of the tetrahedra in the enclosing triangulation. It follows from the above remarks that the two resulting groups of tetrahedra are disconnected from each other in the enclosing triangulation.

Similarly, if the enclosing triangulation is not connected then we can partition the tetrahedra into two groups each of which is disconnected from the other. This partition of tetrahedra then induces a partition of the labels, which from the above remarks leads to a partition of the signature cycles for which no label appears in both groups of cycles. ■

Example 4.4.10 Signature $(abeb)(ccd)(ea)(d)$ does not have a connected enclosing triangulation since its cycles can be partitioned as $(abeb)(ea) \mid (ccd)(d)$. Note that none of the five labels appear on both sides of the partition.

Lemma 4.4.11 The underlying surface of a signature is two-sided in the enclosing triangulation if and only if the cycles that form the signature can be partitioned into two groups so that every label appears in both groups.

Proof Within each tetrahedron of the enclosing triangulation, consider the left and right arrows to be lifted slightly away from the corresponding quadrilateral. In particular, each arrow is lifted in the direction of its parallel hinge edge as illustrated in Figure 4.12. Note that within each path of left and right arrows, the tip of each arrow still meets the following arrow at its base, even after this lifting has been performed. That is, adjacent arrows within a path are lifted away from the splitting surface in the same direction. This can be seen by examining all possible edge identifications as described in Lemma 4.2.3 and following the corresponding tetrahedron face identification in each case.




Figure 4.12: Lifting left and right arrows away from a quadrilateral

After this lifting is performed, each path within a natural representation of the underlying surface becomes a path along one particular side of the surface. If the underlying surface is two-sided within the enclosing triangulation, these paths can be partitioned into two groups according to which side of the surface they lie on. This then induces a partition of the signature cycles. Furthermore, each quadrilateral label q appears in both groups of this partition of cycles since the corresponding arrows q^- and q^+ lie on opposite sides of the underlying surface.

On the other hand, assume that the underlying surface is one-sided within the enclosing triangulation. Then there is some closed curve ρ within a neighbourhood of the underlying surface that cuts through this surface an odd number of times. Let ρ' be the projection of this curve onto the surface itself; we can manipulate ρ' so that it does not meet any vertices of the surface as illustrated in the left hand diagram of Figure 4.13.




Figure 4.13: A fragment of curves ρ' and ρ''

We see then that we can create a new closed curve ρ'' that follows alongside ρ' , where the curve ρ'' is constructed by following our lifted arrows backwards and forwards along paths and occasionally crossing through some quadrilateral q from arrow q^- to q^+ or vice versa. Such a curve ρ'' is illustrated in the right hand diagram of Figure 4.13. Note that ρ'' always lies on the same side of the surface as the arrows that it follows; in particular ρ'' might not always lie on the same side of the surface as ρ . Nevertheless, since ρ crosses the surface an odd number of times we see that ρ'' must also cross the surface an odd number of times.

Assume then that the cycles of our signature σ can be partitioned into two groups so that every quadrilateral label appears in both groups. Recall that each appearance of a quadrilateral label q within σ corresponds to an arrow q^- or q^+ within the natural representation of the underlying surface. We can thus track the closed curve ρ'' within our signature. Each step within ρ'' that follows an arrow forwards or backwards corresponds to moving one step forwards or backwards within a cycle of σ , and each step within ρ'' that cuts through the surface corresponds to jumping between a left arrow q^- and its corresponding right arrow q^+ or vice versa.

Observe that each step forwards or backwards within a cycle leaves us within the same group of the cycle partition, and that each jump between corresponding left and right arrows switches between groups of this partition. Because the curve ρ'' is closed, this sequence of steps must end where it begins and so we must switch between groups an even number of times. This contradicts the earlier observation that the path ρ' cuts through the surface an odd number of times. Therefore we cannot partition the cycles of σ as initially claimed. ■

Example 4.4.12 The underlying surface of signature $(ab)(bc)(a)(c)$ is two-sided within the enclosing triangulation since its cycles can be partitioned as $(ab)(c) \mid (a)(bc)$. Note that all three labels appear on both sides of the partition.

The underlying surface of signature $(abb)(ac)(c)$ is one-sided within the enclosing triangulation since for any partition of its cycles the two instances of label b must appear on the same side of the partition.

Definitions 4.4.13 A signature is *connected* if its enclosing triangulation is connected; otherwise the signature is *disconnected*. A signature is *two-sided* if its underlying surface is two-sided in the enclosing triangulation; otherwise the signature is *one-sided*.

We begin now to investigate combinatorial properties of the enclosing triangulation. In particular we are interested in the numbers of vertices and edges of this triangulation and whether this triangulation is in fact a 3-manifold triangulation. For these calculations we must examine equivalence classes of signature positions as follows.

Definition 4.4.14 (Signature Positions) Consider a signature σ of order n . If we ignore brackets, σ forms a sequence of $2n$ labels when written from left to right, where each label appears precisely twice. A *signature position* is simply an integer between 1 and $2n$ inclusive representing a position within this sequence of labels. The set of all signature positions in σ is denoted by $\mathcal{P}(\sigma)$, i.e., $\mathcal{P}(\sigma) = \{1, \dots, 2n\}$.

Notation For any label q within a signature, we let q_1 denote the signature position corresponding to the first occurrence of q and we let q_2 denote the signature position corresponding to the second occurrence of q (thus $q_1 < q_2$).

For instance, in the signature $(ab)(bc)(a)(c)$ the first occurrence of label a is at position 1 and the second occurrence is at position 5. Thus $a_1 = 1$ and $a_2 = 5$.

Definitions 4.4.15 Let σ be a signature. The *cycle operation* c and the *swap operation* s are permutations on $\mathcal{P}(\sigma)$ defined as follows. Note that we use right function notation with the cycle and swap operations, i.e., we write xc and xs instead of $c(x)$ and $s(x)$.

- For signature position x , if x is not at the end of a cycle then $xc = x + 1$. Otherwise xc is the position corresponding to the beginning of the cycle containing x . Thus the cycle operation has the effect of moving to the next position within a cycle.
- For any label q , if x is the signature position of one occurrence of q then xs is the signature position of the other occurrence of q . That is, for any label q we define $q_1s = q_2$ and $q_2s = q_1$. Note that s^2 is simply the identity permutation.

Definitions 4.4.16 Let σ be a signature. We define $h(\sigma)$ to be the number of hinge edges in the enclosing triangulation, $v^2(\sigma)$ to be the number of vertices in the underlying surface, $v^3(\sigma)$ to be the number of vertices in the enclosing triangulation, $\chi^2(\sigma)$ to be the Euler characteristic of the underlying surface and $\chi^3(\sigma)$ to be the Euler characteristic of the enclosing triangulation.

Theorem 4.4.17 The quantities defined above can be calculated for any signature σ of order n as follows.

- $h(\sigma)$ is the number of cycles in σ ;
- $v^2(\sigma)$ is the number of equivalence classes defined by permutations s and $cscsc^{-1}sc^{-1}s$, i.e., if we consider two signature positions to be equivalent when they are related by any combinations of permutations s , $cscsc^{-1}sc^{-1}s$ and their inverses then $v^2(\sigma)$ is the number of equivalence classes that result;
- $v^3(\sigma)$ is the number of equivalence classes defined by permutations scs and $cscsc^{-1}$;
- $\chi^2(\sigma) = v^2(\sigma) - n$;
- $\chi^3(\sigma) = v^3(\sigma) - h(\sigma) - \chi^2(\sigma)$.

Proof Consider the cycle and swap operations c and s . Each of these operations transforms one signature position into another. By considering the left and right arrows within a natural representation that correspond to these signature positions, we can extend c and s to operations that transform one arrow into another.

In this context, the cycle operation c transforms an arrow α into the subsequent arrow along the path containing α . The swap operation s transforms a left arrow q^- into the corresponding right arrow q^+ and vice versa.

Armed with this knowledge, we now prove each claim in turn.

- Recall from the proof of Theorem 4.1.4 that each hinge edge in the enclosing triangulation is surrounded by a ring of prisms, as illustrated in the left hand diagram of Figure 4.14. The shaded quadrilaterals in this diagram represent the quadrilaterals of the underlying splitting surface.




Figure 4.14: Quadrilaterals surrounding a hinge edge

We observe then that each hinge edge is surrounded by a single path of arrows, as illustrated by the dotted arrows in the right hand diagram of Figure 4.14. In this way the hinge edges of the enclosing triangulation correspond precisely to the paths of arrows, which in turn correspond precisely to the cycles within the signature σ . Thus $h(\sigma)$ is the number of cycles in σ .

- Let V be a vertex of the underlying surface. The link of vertex V is a disc formed from some number of quadrilateral corners.

Recall Lemma 4.3.1 which states that each upper edge of a quadrilateral in the underlying surface is identified with some lower edge of a quadrilateral. Combining this result with a study of the eight oriented edge identifications illustrated in Figure 4.8 on page 134, we see that the link of V can in fact be represented as a sequence of quadrilateral corners c_0, c_1, \dots, c_{4k} for some k in which the following condition is satisfied. Each corner c_{4i} is at the top of a quadrilateral, each corner c_{4i+2} is at the base of a quadrilateral and each of the remaining $2k$ corners is on the side of a quadrilateral. A fragment of such a vertex link is illustrated in the left hand diagram of Figure 4.15, where each quadrilateral is displayed upright according to its diamond representation.




Figure 4.15: A fragment of a vertex link in the underlying surface

In particular, it follows that each vertex is found at the top of some quadrilateral. For each left or right arrow α , let V_α denote the vertex at the top of the quadrilateral containing α . We then observe for any pair of arrows α and β that vertices V_α and V_β are identified if and only if arrows α and β can be related by some combination of the following transformations.

- A left arrow q^- can be converted to the corresponding right arrow q^+ or vice versa. This transformation is represented by the operation s .
- An arrow α can be transformed by circling 360° around the vertex V_α . This is illustrated in the right hand diagram of Figure 4.15 in which arrow α is transformed into arrow β . This transformation is represented by the sequence of operations $cscsc^{-1}sc^{-1}s$.

It follows then that the number of vertices of the underlying surface is the number of equivalence classes of arrows defined by the two transformations listed above. Translating this result back into the notation of signature positions we see that $v^2(\sigma)$ is the number of equivalence classes of signature positions defined by the signature position permutations s and $cscsc^{-1}sc^{-1}s$.

- Consider any arrow α in the underlying surface. If Δ is the tetrahedron containing this arrow in the enclosing triangulation, we let V_α denote the vertex of Δ that is pointed to by α . Note that α does not actually pass through V_α since V_α is a vertex of the enclosing triangulation and does not lie on the underlying surface. Figure 4.16 illustrates the vertex V_α for a single arrow α .

It can be observed that precisely two vertices of each tetrahedron are pointed to by an arrow within this tetrahedron. However, since every tip of an arrow meets the base of some other arrow, it follows that the remaining two vertices of each tetrahedron are pointed to by an




Figure 4.16: The tetrahedron vertex pointed to by an arrow

arrow within an adjacent tetrahedron. Thus each vertex of the enclosing triangulation can be written as V_α for some arrow α .

Consider now the ways in which the different tetrahedron vertices are identified in the enclosing triangulation. It can be shown for any pair of arrows α and β that vertices V_α and V_β are identified in the enclosing triangulation if and only if arrows α and β can be related by some combination of the following transformations.

- An arrow can be converted into a parallel arrow in an adjacent quadrilateral. This is illustrated in the left hand diagram of Figure 4.17 in which arrow α is transformed into arrow β . This transformation is represented by the sequence of operations scs .




Figure 4.17: Transformations involving parallel arrows

- An arrow α placed behind some arrow α' can be converted into the arrow β which is behind arrow β' , where arrows α' and β' are parallel arrows in adjacent quadrilaterals. This transformation is illustrated in the right hand diagram of Figure 4.17 and is represented by the sequence of operations $cscsc^{-1}$.

We see then that the number of vertices in the enclosing triangulation is the number of equivalence classes of arrows defined by these two transformations. Translating this result into the language of signature positions we see that $v^3(\sigma)$ is the number of equivalence classes of signature positions defined by the signature position permutations scs and $cscsc^{-1}$.

- The quadrilateral structure of the underlying surface contains precisely n quadrilateral faces, $2n$ edges and $v^2(\sigma)$ vertices. Thus $\chi^2(\sigma) = v^2(\sigma) - 2n + n = v^2(\sigma) - n$.
- Consider the enclosing triangulation. This triangulation contains precisely n tetrahedra, $2n$ triangular faces and $v^3(\sigma)$ vertices. Furthermore, its edges are either hinge edges or edges cutting through the underlying surface.

There are precisely $h(\sigma)$ hinge edges and $v^2(\sigma)$ remaining edges, since the edges cutting through the underlying surface meet this surface in its vertices and each vertex of the un-

derlying surface lies on such an edge. Therefore $\chi^3(\sigma) = v^3(\sigma) - h(\sigma) - v^2(\sigma) + 2n - n = v^3(\sigma) - h(\sigma) - \chi^2(\sigma)$.

This concludes our series of proofs. ■

Example 4.4.18 Consider the signature $\sigma = (abc)(a)(cb)$, in which positions 1, 2, 3, 4, 5 and 6 are represented by symbols a_1, b_1, c_1, a_2, c_2 and b_2 respectively. We can use Theorem 4.4.17 to calculate the quantities $h(\sigma)$, $v^2(\sigma)$, $v^3(\sigma)$, $\chi^2(\sigma)$ and $\chi^3(\sigma)$.

We begin by observing that $h(\sigma) = 3$ since σ contains three cycles. Observe that the swap operation s is simply the permutation $(a_1a_2)(b_1b_2)(c_1c_2)$ and that the cycle operation is the permutation $(a_1b_1c_1)(a_2)(c_2b_2)$, using the usual cycle notation for permutations.

Permutation $cscsc^{-1}sc^{-1}s$ is then $(a_1c_1b_1)(a_2b_2c_2)$. We thus see that any two signature positions can be related by some combination of permutations s and $cscsc^{-1}sc^{-1}s$, and so $v^2(\sigma) = 1$.

Permutations scs and $cscsc^{-1}$ can be calculated as $(a_1)(b_1c_1)(a_2b_2c_2)$ and $(a_1b_1)(c_1)(a_2c_2b_2)$ respectively. Here we see that each of a_1, b_1 and c_1 can be related through a combination of these two permutations and similarly that each of a_2, b_2 and c_2 can be related through a combination of these permutations. Positions a_1 and a_2 however cannot be related using these permutations and so permutations scs and $cscsc^{-1}$ divide the signature positions into the equivalence classes $\{a_1, b_1, c_1\}$ and $\{a_2, b_2, c_2\}$. Thus $v^3(\sigma) = 2$.

Finally we can calculate $\chi^2(\sigma) = v^2(\sigma) - n = 1 - 3 = -2$ and $\chi^3(\sigma) = v^3(\sigma) - h(\sigma) - \chi^2(\sigma) = 2 - 3 + 2 = 1$.

Corollary 4.4.19 Let σ be a connected signature of order n . Then the enclosing triangulation of σ is a 3-manifold triangulation if and only if $v^3(\sigma) - h(\sigma) = v^2(\sigma) - n$. Otherwise it is true that $v^3(\sigma) - h(\sigma) > v^2(\sigma) - n$.

Proof Let T be the enclosing triangulation of σ . Recall from Lemma 4.4.5 that no edge of triangulation T is identified with itself in reverse. Therefore T is a 3-manifold triangulation precisely when $\chi(T) = 0$, and otherwise it is known that $\chi(T)$ is strictly positive.

The result then follows from Theorem 4.4.17 in which it is proven that $\chi(T) = \chi^3(\sigma) = v^3(\sigma) - h(\sigma) - \chi^2(\sigma) = v^3(\sigma) - h(\sigma) - v^2(\sigma) + n$. ■

Lemma 4.4.20 For any signature σ of order n , it is true that $v^2(\sigma) \leq n$, that $v^3(\sigma) \leq h(\sigma)$ and that $\chi^2(\sigma) \leq 0$.

Proof We examine each inequality in turn.

- Since permutation s contains precisely n cycles, there are n equivalence classes defined by permutation s . Thus there can be at most n equivalence classes defined by any collection of permutations that includes s , and so $v^2(\sigma) \leq n$.
- Since $s = s^{-1}$, permutation scs is a conjugate of c and thus contains the same number of cycles as c . This in turn is the number of cycles in σ which is $h(\sigma)$. Thus there are precisely $h(\sigma)$ equivalence classes defined by permutation scs , and so there can be at most $h(\sigma)$ equivalence classes defined by any collection of permutations that includes scs . Hence $v^3(\sigma) \leq h(\sigma)$.

- In Theorem 4.4.17 it is established that $\chi^2(\sigma) = v^2(\sigma) - n$. From an earlier result of this lemma we see that $v^2(\sigma) \leq n$, and so $\chi^2(\sigma) \leq 0$.

This concludes our series of proofs. ■

Lemma 4.4.21 If σ is a connected signature of order n then $v^3(\sigma) \leq h(\sigma) \leq n + 1$.

Proof Recall from the proof of Lemma 4.4.9 that two tetrahedra Δ and Δ' are adjacent in the enclosing triangulation of σ if and only if their corresponding quadrilateral labels q and q' are adjacent symbols within some cycle of σ . Since σ and therefore its enclosing triangulation are connected, it follows that any two signature positions can be related through a combination of cycle operations and swap operations.

Therefore, if permutation c contains n_c cycles of lengths c_1, \dots, c_{n_c} and permutation s contains n_s cycles of lengths s_1, \dots, s_{n_s} , we see that $\sum(c_i - 1) + \sum(s_j - 1) \geq 2n - 1$, since there are $2n$ signature positions all of which must be reachable from each other using operations c and s .

But $\sum(c_i - 1) + \sum(s_j - 1) = (2n - n_c) + (2n - n_s)$, and so we see that $n_c + n_s \leq 2n + 1$. Since s contains precisely n cycles each containing two elements, we have $n_s = n$ and therefore $n_c \leq n + 1$.

We finally observe that the cycles of c correspond precisely to the cycles of σ and so $n_c = h(\sigma)$, from which we conclude that $h(\sigma) \leq n + 1$. Using Lemma 4.4.20 we can extend this result to $v^3(\sigma) \leq h(\sigma) \leq n + 1$. ■

4.5 Census of Signatures

As we approach the conclusion of this chapter we present the results of a census of all connected signatures of order ≤ 8 . A census of signatures, like a census of triangulations, provides a wealth of examples from which to draw conjectures and test hypotheses. All calculations for this census were performed using the program *Regina* which is described in detail in Section 1.3.

In this census we consider each signature precisely once up to the operations described in Lemma 4.4.8. That is, two signatures are not both included in the census if they are related by a sequence of these operations.

We begin by examining large-scale statistics of the census in Table 4.1. Following this we present in Section 4.5.1 full tables of all signatures of order ≤ 3 along with various combinatorial properties of these signatures. In Section 4.5.2 we examine in detail the much smaller class of signatures whose enclosing triangulations are true 3-manifold triangulations.

Recall that Chapter 3 presents a list of all prime minimal 3-manifold triangulations formed from ≤ 6 tetrahedra. Although this list does not extend to larger numbers of tetrahedra it has been possible to establish precisely which signatures of orders 7 and 8 have prime minimal enclosing triangulations. The techniques of Section 2.4 have been used to discard triangulations that can be identified as non-prime or non-minimal, and the remaining triangulations have all been identified as members of the families described in Section 3.3. Their underlying 3-manifolds have thus been identified, and their minimality has been verified through the computational results of Martelli and Petronio [26].

A variety of statistics for the census of signatures is presented in Table 4.1. Each row of this table represents the set of all connected signatures of a particular order. The individual table columns have the following meanings.

- *Order*: The order n of the signatures described by this row of the table.
- *Signatures*: The total number of connected signatures of order n .
- *3-Manifold*: The number of connected signatures of order n whose enclosing triangulations are in fact 3-manifold triangulations.
- *1-Vertex*: The number of connected signatures of order n whose enclosing triangulations are one-vertex 3-manifold triangulations.
- *Prime Minimal*: The number of connected signatures of order n whose enclosing triangulations are prime minimal 3-manifold triangulations.

Recall from Corollary 2.4.9 that all closed orientable prime minimal triangulations with ≥ 3 tetrahedra have only one vertex. For triangulations with ≤ 2 tetrahedra however this is not necessarily true, as is seen in the first two rows of the table.

- *Time*: The running time taken to calculate all connected signatures of order n . Running times are measured on a single 1.2GHz Pentium III processor and are displayed as $h:mm:ss$ unless otherwise indicated.

Order	Signatures	3-Manifold	1-Vertex	Prime Minimal	Time
1	2	2	1	2	0:00
2	5	5	3	4	0:00
3	20	12	7	3	0:00
4	95	36	20	3	0:00
5	636	105	61	5	0:00
6	5 731	452	280	8	0:05
7	66 540	1 759	1 029	10	5:51
8	952 965	8 942	5 246	16	10:23:08

Table 4.1: Statistics for the connected signature census

4.5.1 Small Signatures in Detail

In Tables 4.2, 4.3 and 4.4 we enumerate all connected signatures of orders 1, 2 and 3 respectively and present several of their combinatorial properties as defined in Definitions 4.4.16. The connected signatures of larger orders are too numerous to list here although the results are available from the *Regina* website [3].

Signatures σ are arranged in these tables according to $v^3(\sigma)$, i.e., the number of vertices in the enclosing triangulation, and $v^2(\sigma)$, i.e., the number of vertices in the underlying surface. The number of hinge edges $h(\sigma)$ can be seen directly from the presentation of each signature since Theorem 4.4.17 shows this to be the number of cycles in σ . Recall from Theorem 4.4.17 that the Euler characteristics of the underlying surface and the enclosing triangulation can be calculated from these quantities as $\chi^2(\sigma) = v^2(\sigma) - n$ and $\chi^3(\sigma) = v^3(\sigma) - h(\sigma) - \chi^2(\sigma)$.

	$v^2(\sigma) = 1$
$v^3(\sigma) = 1$	(aa)
$v^3(\sigma) = 2$	$(a)(a)$

Table 4.2: All connected signatures of order 1

	$v^2(\sigma) = 1$	$v^2(\sigma) = 2$
$v^3(\sigma) = 1$	$(aab)(b)$	$(aabb)$ $(abab)$
$v^3(\sigma) = 2$		$(ab)(ab)$
$v^3(\sigma) = 3$		$(ab)(a)(b)$

Table 4.3: All connected signatures of order 2

4.5.2 Signatures with 3-Manifold Triangulations

It is seen in Table 4.1 that an unhappily small number of signatures have enclosing triangulations that are true 3-manifold triangulations. The 55 such signatures of order ≤ 4 and the 3-manifolds represented by their enclosing triangulations are presented in Table 4.5. The signatures of larger orders with enclosing 3-manifold triangulations are too numerous to present here.

We turn now to the set of all signatures of order ≤ 8 whose enclosing triangulations are not only 3-manifold triangulations but also prime and minimal. This set is so small that we can present all 51 of these signatures and their corresponding 3-manifolds in Table 4.6. Signatures in this table whose enclosing triangulations have more than one vertex are marked as such.

Note that the only splitting surfaces that we find within closed orientable prime minimal 3-manifold triangulations of ≥ 3 tetrahedra are in fact non-orientable splitting surfaces. This can be explained by Corollary 2.4.9, which states that any such triangulation must have only one vertex. As a result, any splitting surface within such a triangulation must be one-sided and therefore non-orientable.

4.6 Future Directions

As discussed in opening comments of this chapter, research into splitting surfaces is still in progress. In particular, planned directions of research include the following.

- Recall that the results presented within this chapter from Section 4.3 onwards pertain only to splitting surfaces that allow for natural representations. As seen in Example 4.2.8, there are splitting surfaces that do not allow for such representations. For instance, splitting surfaces with no natural representations are seen within the two-tetrahedron and three-tetrahedron minimal triangulations of the orientable spaces $L(6,1)$ and $L(10,3)$ as well as the non-orientable spaces $S^2 \times S^1$ and $\mathbb{R}P^2 \times S^1$.

It is desirable then to expand the concept of a signature to allow for descriptions of non-natural representations of splitting surfaces. This can be achieved by allowing left arrows

	$v^2(\sigma) = 1$	$v^2(\sigma) = 2$	$v^2(\sigma) = 3$
$v^3(\sigma) = 1$	$(aabbcc)$ $(abcabc)$ $(aabb)(c)$ $(ababc)(c)$ $(aabc)(bc)$ $(abac)(bc)$ $(abac)(b)(c)$ $(aab)(bc)(c)$ $(ab)(ac)(bc)$	$(aabc)(c)$ $(abcacb)$ $(aab)(bcc)$	$(abcabc)$ $(abccba)$
$v^3(\sigma) = 2$	$(abc)(ab)(c)$ $(ab)(ac)(b)(c)$	$(aabc)(b)(c)$	$(abc)(abc)$ $(abc)(acb)$
$v^3(\sigma) = 3$			
$v^3(\sigma) = 4$			$(abc)(a)(b)(c)$

Table 4.4: All connected signatures of order 3

and right arrows within a path to be traversed in either the forward or the reverse direction. If we declare the symbol \bar{q} within a signature to represent an arrow traversed in the reverse direction, we can for instance represent the non-natural diamond representation of Figure 4.7 (seen on page 133) with the signature $(aa\bar{c})(bbc)$.

Not only will such a generalisation of signatures allow us to represent any splitting surface within a 3-manifold triangulation, but it will also allow us to represent splitting surfaces within non-orientable triangulations which thus far we have been unable to do. Signatures of this more complex form are however more difficult to analyse, and research into their properties is continuing.

- Recall Section 3.3, in which we present a small number of infinite parameterised families of triangulations that together include almost every closed orientable prime minimal triangulation formed from ≤ 6 tetrahedra. In a similar fashion we can create infinite parameterised families of signatures. By forming families of signatures in this way, the properties of the corresponding underlying surfaces and enclosing triangulations can be more easily analysed.

The existence of such families is already apparent from Table 4.6. Every signature of the form $(a_1a_2 \dots a_na_n \dots a_2a_1)$ can be shown to have an enclosing triangulation representing the lens space $L(4n, 2n - 1)$. In addition every signature of the form $(a_1a_2 \dots a_na_1a_2 \dots a_n)$ can be shown to have an enclosing triangulation representing the quotient space S^3/Q_{4n} (this triangulation is in fact a twisted layered loop as described in Section 3.3.2). We aim to identify and analyse more such families.

- It can be seen from Table 4.1 that there are vastly more signatures than there are signatures whose enclosing triangulations are in fact 3-manifold triangulations. For this reason it is desirable to expand upon the results of Section 4.4 and obtain further methods of extracting information from a signature regarding its enclosing triangulation.

Signature	3-Man.	Signature	3-Man.	Signature	3-Man.
(aa)	$L(4, 1)$	(aabcdlcb)	$L(16, 7)$	(abac)(bdd)(c)	$L(4, 1)$
(a)(a)	S^3	(abcdabcd)	S^3/Q_{16}	(abac)(bd)(cd)	$\mathbb{R}P^3$
(aabb)	$L(8, 3)$	(aabcdcb)(d)	$\mathbb{R}P^3$	(abac)(bd)(c)(d)	$\mathbb{R}P^3$
(abab)	S^3/Q_8	(aabccd)(b)(d)	$\mathbb{R}P^3 \# \mathbb{R}P^3$	(abcd)(abcd)	$L(4, 1)$
(aab)(b)	$\mathbb{R}P^3$	(aabcdb)(c)(d)	$\mathbb{R}P^3$	(abcd)(adcb)	$L(4, 1)$
(ab)(ab)	$\mathbb{R}P^3$	(abacbd)(cd)	$\mathbb{R}P^3 \# \mathbb{R}P^3$	(abcd)(ac)(bd)	$\mathbb{R}P^3$
(ab)(a)(b)	S^3	(abacdc)(b)(d)	$S^2 \times S^1$	(abcd)(a)(b)(c)(d)	S^3
(aabccb)	$L(12, 5)$	(abcabd)(c)(d)	$S^2 \times S^1$	(aab)(bcd)(c)(d)	$L(4, 1)$
(abcabc)	S^3/Q_{12}	(aabbcc)(cd)(d)	$L(8, 3)$	(aab)(bc)(cd)(d)	$\mathbb{R}P^3$
(aabcb)(c)	$\mathbb{R}P^3$	(aabcb)(cdd)	$L(14, 3)$	(abc)(abd)(cd)	$L(8, 3)$
(aabc)(b)(c)	$\mathbb{R}P^3$	(aabcb)(cd)(d)	$L(8, 3)$	(abc)(abd)(c)(d)	S^3
(abac)(b)(c)	$S^2 \times S^1$	(aabcd)(b)(c)(d)	$\mathbb{R}P^3$	(abc)(acd)(bd)	S^3/Q_8
(aab)(bcc)	$L(6, 1)$	(ababc)(cd)(d)	S^3/Q_8	(abc)(acd)(b)(d)	S^3
(aab)(bc)(c)	$L(4, 1)$	(abacd)(b)(c)(d)	$S^2 \times S^1$	(abc)(ab)(cd)(d)	$\mathbb{R}P^3$
(abc)(abc)	$L(3, 1)$	(aabc)(bd)(cd)	$S^2 \times S^1$	(abc)(ad)(bd)(c)	$\mathbb{R}P^3$
(abc)(acb)	$L(3, 1)$	(aabc)(bd)(c)(d)	$\mathbb{R}P^3$	(abc)(ad)(b)(c)(d)	S^3
(abc)(a)(b)(c)	S^3	(abac)(bcdcd)	$L(4, 1)$	(ab)(ac)(bd)(cd)	S^3
(ab)(ac)(bc)	$L(4, 1)$	(abac)(bcd)(d)	$L(4, 1)$	(ab)(ac)(bd)(c)(d)	S^3
(ab)(ac)(b)(c)	S^3				

Table 4.5: All signatures of order ≤ 4 with enclosing 3-manifold triangulations

In particular, in a fashion similar to the results of Section 2.3.1, we aspire to prove results allowing us whilst constructing a signature to identify that its enclosing triangulation cannot possibly be a 3-manifold triangulation. Such results can then be incorporated into the signature census algorithm, allowing us to prune the recursive construction of signatures and thereby improve the efficiency of the algorithm.

Signature	3-Manifold	Signature	3-Manifold
(aa)	$L(4, 1)$	(abcdefgabcdefg)	S^3/Q_{28}
(a)(a)	S^3 (2 vertices)	(aabcdefedcb)(fgg)	$L(38, 9)$
(aabb)	$L(8, 3)$	(aabcdedcb)(efggf)	$L(62, 23)$
(abab)	S^3/Q_8	(aabcdec)(dfegfg)	$S^3/Q_{48} \times \mathbb{Z}_5$
(aab)(b)	$\mathbb{R}P^3$	(abacdedc)(bfff)(egg)	$L(40, 9)$
(ab)(ab)	$\mathbb{R}P^3$ (2 vertices)	(aabcdcb)(defggfe)	$L(70, 29)$
(aabcbb)	$L(12, 5)$	(aabcdcb)(defe)(fgg)	$L(32, 5)$
(abcabc)	S^3/Q_{12}	(abacdc)(beffe)(dgg)	$L(52, 11)$
(aab)(bcc)	$L(6, 1)$	(abac)(bded)(cff)(egg)	$L(10, 1)$
(aabcdccb)	$L(16, 7)$	(aabcdefghg fedcb)	$L(32, 15)$
(abcdabcd)	S^3/Q_{16}	(abcdefghabcdefgh)	S^3/Q_{32}
(aabc)(cdd)	$L(14, 3)$	(aabcdefgfedcb)(ghh)	$L(46, 11)$
(aabcdeedcb)	$L(20, 9)$	(aabcdefedcb)(fghhg)	$L(78, 29)$
(abcdeabcde)	S^3/Q_{20}	(aabcdefdcb)(egfhhg)	$S^3/Q_{64} \times \mathbb{Z}_7$
(aabcdcb)(dee)	$L(22, 5)$	(abacdefedc)(bgg)(fhh)	$L(56, 13)$
(aabc)(cdeed)	$L(30, 11)$	(aabcdedcb)(efghhg)	$L(94, 39)$
(abac)(bdd)(cee)	$L(8, 1)$	(aabcdedcb)(efgf)(ghh)	$L(44, 7)$
(aabcdef fedcb)	$L(24, 11)$	(abacdedc)(bfggf)(ehh)	$L(84, 19)$
(abcdefabcdef)	S^3/Q_{24}	(aabcdcb)(defgfe)(ghh)	$L(80, 17)$
(aabcdedcb)(eff)	$L(30, 7)$	(abacdec)(dfegfg)(bhh)	$S^3/Q_{56} \times \mathbb{Z}_5$
(aabcdcb)(deffe)	$L(46, 17)$	(abacbd)(cdefe)(fghhg)	$S^3/Q_{56} \times \mathbb{Z}_3$
(aabcdb)(cedfef)	$S^3/Q_{32} \times \mathbb{Z}_3$	(abacdc)(beffe)(dghhg)	$L(112, 41)$
(abacdc)(bee)(dff)	$L(24, 5)$	(abacdc)(beffe)(dgg)(fhh)	$L(34, 5)$
(aabc)(cded)(eff)	$L(20, 3)$	(abacde)(cdefgf)(bhgh)	SFS ($S^2 : (2, 1),$ (4, 1), (6, -5))
(abac)(bded)(cfef)	S^3/Q_{24}		
(aabcdefgg fedcb)	$L(28, 13)$	(aabc)(cded)(efgf)(ghh)	$L(26, 3)$

Table 4.6: All signatures of order ≤ 8 with enclosing prime minimal 3-manifold triangulations

Chapter 5

Bounding the Projective Solution Space

As noted in Section 1.1, the process of enumerating all vertex normal surfaces for a given triangulation is computationally expensive. It is thus of interest to establish theoretical bounds upon various aspects of this process, including the time required for the computation and the size of the projective solution space.

In this and the following chapter we aim to bound the number of maximal embedded faces of the projective solution space. By examining this as an algebraic problem in edge weight space, we almost produce a bound that is sharp, at least as far as the algebraic problem is concerned. This bound is described in Consequence 5.5.15, and in Section 6.5 we see where the attempted proof breaks down. Although the final result remains unproven, it is the author's belief that the methods used are of interest and that it is possible to adapt these methods to work around the difficulties of Section 6.5.

Chapter 5 deals with the structure of the projective solution space and in particular its realisation in edge weight space. Chapter 6 attempts to pin down the conjectured bound when the pieces used to build the projective solution space are placed in general position.

Before we proceed, recall the following facts from Section 1.1. Normal surfaces in an n -tetrahedron triangulation can be represented as vectors in \mathbb{R}^{7n} with non-negative integer coordinates. Each tetrahedron has seven associated coordinates in this vector. Specifically there are four coordinates for the different types of triangular normal disc and three coordinates for the different types of quadrilateral normal disc within the tetrahedron.

The projective solution space is the entire solution space of the matching equations projected onto the hyperplane in which all vector coordinates sum to 1. This projection is done by a simple scaling, i.e., by dividing a (non-zero) normal surface vector by the sum of its coordinates. This projective solution space is then a finite convex polytope embedded in \mathbb{R}^{7n} with all points of the polytope having non-negative coordinates.

Recall also that a normal surface vector corresponds to an embedded normal surface if and only if at most one quadrilateral coordinate associated with each tetrahedron is non-zero.

Assumptions Throughout this chapter every triangulation is assumed to be a one-vertex triangulation with n tetrahedra.

5.1 Pseudo-Convexity and Faces

Before we continue it is necessary to make some definitions regarding convexity and faces of polytopes and cell complexes. Although the cell complexes that we use throughout Chapters 5 and 6 are not always convex, they are also not entirely arbitrary. Recall that a convex set \mathcal{C} is one in which for any distinct points $x, y \in \mathcal{C}$ the entire closed interval $[x, y]$ lies within \mathcal{C} . It is proven later in this chapter that the spaces with which we work satisfy the slightly broader property of pseudo-convexity as defined below.

Definition 5.1.1 (Pseudo-Convexity) Set \mathcal{C} is defined to be *pseudo-convex* if for any convex subset $\mathcal{S} \subseteq \mathcal{C}$ and for any distinct points $x, y \in \mathcal{C}$ the following condition is satisfied. If the open interval (x, y) contains some point in \mathcal{S} , then there is some convex subset $\mathcal{S}' \subseteq \mathcal{C}$ containing both \mathcal{S} and the entire closed interval $[x, y]$, i.e., $(\mathcal{S} \cup [x, y]) \subseteq \mathcal{S}'$.




Figure 5.1: A pseudo-convex and non-pseudo-convex set in \mathbb{R}^2

Example 5.1.2 Figure 5.1 illustrates two subsets of the real plane. The figure on the left formed from three line segments meeting at a common vertex is pseudo-convex. The figure on the right formed from a square and a line segment meeting at a common vertex is not pseudo-convex, since if set \mathcal{S} represents the continuation of the line through the square and points x and y are as marked then the property required by Definition 5.1.1 is not satisfied.




Figure 5.2: A pseudo-convex and non-pseudo-convex set in \mathbb{R}^3

In Figure 5.2 we see two subsets of real 3-space. The left figure is formed from a horizontal triangle and a vertical line segment meeting at a common vertex, and the right figure is formed from a horizontal square and a vertical line passing through each others' interiors. It should be noted that in each diagram of Figure 5.2 the line segment and the shaded polygon do not belong to a common plane. In particular, in most cases a line joining a point of the illustrated line segment with a point of the shaded polygon will not meet any other part of the figure.

Of the two subsets of \mathbb{R}^3 illustrated in Figure 5.2, we once more find that the left hand subset is pseudo-convex whereas the right hand subset is not. In the right hand subset, if we let \mathcal{S} represent the square and x and y represent points of the line above and below the square respectively we can see where the property of Definition 5.1.1 fails to hold.

Some of the similarities between convexity and pseudo-convexity are illustrated in the following results.

Lemma 5.1.3 The following properties are true of pseudo-convexity.

1. Any convex set is also pseudo-convex.
2. The intersection of two pseudo-convex sets is pseudo-convex.

Proof We prove each statement in turn.

1. Let \mathcal{C} be a convex set. Then if we have convex subset $\mathcal{S} \subseteq \mathcal{C}$ meeting open interval (x, y) for $x, y \in \mathcal{C}$, the convexity of \mathcal{C} implies that $[x, y] \subseteq \mathcal{C}$ and so the entire convex set \mathcal{C} satisfies the required property $(\mathcal{S} \cup [x, y]) \subseteq \mathcal{C}$.
2. Let \mathcal{C}_1 and \mathcal{C}_2 be pseudo-convex sets and let $\mathcal{S} \subseteq \mathcal{C}_1 \cap \mathcal{C}_2$ be some convex subset of their intersection meeting open interval (x, y) for $x, y \in \mathcal{C}_1 \cap \mathcal{C}_2$. By pseudo-convexity of \mathcal{C}_1 and \mathcal{C}_2 we can find convex subsets $\mathcal{S}'_1 \subseteq \mathcal{C}_1$ and $\mathcal{S}'_2 \subseteq \mathcal{C}_2$ for which $(\mathcal{S} \cup [x, y]) \subseteq \mathcal{S}'_1$ and $(\mathcal{S} \cup [x, y]) \subseteq \mathcal{S}'_2$. But then $(\mathcal{S} \cup [x, y]) \subseteq \mathcal{S}'_1 \cap \mathcal{S}'_2$, and since \mathcal{S}'_1 and \mathcal{S}'_2 are both convex their intersection is also convex. So $\mathcal{S}' = \mathcal{S}'_1 \cap \mathcal{S}'_2$ is a convex subset of $\mathcal{C}_1 \cap \mathcal{C}_2$ with the required property and hence $\mathcal{C}_1 \cap \mathcal{C}_2$ is pseudo-convex.

This concludes our sequence of proofs. ■

We come now to defining the faces of a pseudo-convex set. Most sources encountered by the author appear to define a face of a convex polytope as the intersection of the polytope with a supporting hyperplane, as seen for instance in [6] and [48]. This definition however does not serve for the non-convex case, even for spaces satisfying the relatively strong condition of pseudo-convexity. We thus use an alternative definition as presented by Brøndsted in [2].

Definition 5.1.4 (Faces) Let \mathcal{C} be a pseudo-convex set. A *face* of \mathcal{C} is a convex subset $\mathcal{F} \subseteq \mathcal{C}$ for which the following condition is satisfied. If x and y are two distinct points in \mathcal{C} for which the open interval (x, y) contains some point in \mathcal{F} , then the entire closed interval $[x, y]$ is in fact contained in \mathcal{F} .

Notice that a face can be of any dimension, possibly ranging all the way from an entire polytope down to individual edges and vertices. Thus, for instance, a tetrahedron has one face of dimension three (the tetrahedron itself), four faces of dimension two (the triangles on its boundary), six faces of dimension one (its edges) and four faces of dimension zero (its vertices).

Brøndsted presents this definition only for convex spaces, and in particular we assume all of the usual results regarding faces of convex polytopes as can be found in most references on the subject such as [2], [6] and [48]. It can be seen furthermore that Definition 5.1.4 coincides with our

usual understanding of what a face should be for a more general pseudo-convex set. It is worth at this point verifying from Definition 5.1.4 some basic properties that we expect faces to satisfy in this larger context.

Lemma 5.1.5 Let \mathcal{C} be a pseudo-convex set. Then the following statements are true.

1. Any maximally convex subset of \mathcal{C} (i.e., a convex subset of \mathcal{C} not contained in any other convex subset of \mathcal{C}) is a face of \mathcal{C} .
2. If \mathcal{F} and \mathcal{G} are both faces of \mathcal{C} then the intersection $\mathcal{F} \cap \mathcal{G}$ is also a face of \mathcal{C} .
3. If \mathcal{F} is a face of \mathcal{C} then any subset $\mathcal{G} \subseteq \mathcal{F}$ is a face of \mathcal{C} if and only if it is a face of \mathcal{F} .

Proof Again we prove each statement in turn.

1. Let \mathcal{F} be a maximally convex subset of \mathcal{C} and suppose that \mathcal{F} meets the open interval (x, y) for some $x, y \in \mathcal{C}$. From the pseudo-convexity of \mathcal{C} there is some convex $\mathcal{S}' \subseteq \mathcal{C}$ containing both \mathcal{F} and $[x, y]$, but by the maximality of \mathcal{F} we must have $\mathcal{S}' = \mathcal{F}$ and thus \mathcal{F} contains all of $[x, y]$. So \mathcal{F} is a face of \mathcal{C} .
2. Observe first that $\mathcal{F} \cap \mathcal{G}$ is an intersection of convex sets and is hence convex. Assume then that $\mathcal{F} \cap \mathcal{G}$ meets open interval (x, y) for some $x, y \in \mathcal{C}$. Then each face \mathcal{F} and \mathcal{G} individually meets (x, y) and so contains all of $[x, y]$. Thus $[x, y] \subseteq \mathcal{F} \cap \mathcal{G}$ and we see that $\mathcal{F} \cap \mathcal{G}$ is a face of \mathcal{C} .
3. Let \mathcal{F} be a face of \mathcal{C} and consider any subset $\mathcal{G} \subseteq \mathcal{F}$. If \mathcal{G} is a face of \mathcal{C} then \mathcal{G} is convex; assume then that \mathcal{G} meets some open interval (x, y) for $x, y \in \mathcal{F}$. Since $\mathcal{F} \subseteq \mathcal{C}$ we have $x, y \in \mathcal{C}$ and so since \mathcal{G} is a face of \mathcal{C} we must have $[x, y] \subseteq \mathcal{G}$. Thus \mathcal{G} is a face of \mathcal{F} .

Suppose on the other hand that \mathcal{G} is a face of \mathcal{F} . Again \mathcal{G} is convex, so assume that \mathcal{G} meets some open interval (x, y) for $x, y \in \mathcal{C}$. Then since $\mathcal{G} \subseteq \mathcal{F}$ we see that \mathcal{F} meets (x, y) and so $[x, y] \subseteq \mathcal{F}$. In particular we have $x, y \in \mathcal{F}$ and so since \mathcal{G} is a face of \mathcal{F} we have $[x, y] \subseteq \mathcal{G}$. Hence \mathcal{G} is a face of \mathcal{C} .

This concludes the series of proofs. ■

Lemma 5.1.6 If \mathcal{F}_1 and \mathcal{F}_2 are faces of pseudo-convex sets \mathcal{C}_1 and \mathcal{C}_2 respectively, then $\mathcal{F}_1 \cap \mathcal{F}_2$ is a face of the pseudo-convex set $\mathcal{C}_1 \cap \mathcal{C}_2$.

Proof Note first that $\mathcal{C}_1 \cap \mathcal{C}_2$ is pseudo-convex from Lemma 5.1.3 and that $\mathcal{F}_1 \cap \mathcal{F}_2$ is an intersection of convex sets and is thus itself convex. If $\mathcal{F}_1 \cap \mathcal{F}_2$ meets some open interval (x, y) for $x, y \in \mathcal{C}_1 \cap \mathcal{C}_2$, then \mathcal{F}_1 and \mathcal{F}_2 each individually meet (x, y) and \mathcal{C}_1 and \mathcal{C}_2 each individually contain both x and y . Thus \mathcal{F}_1 and \mathcal{F}_2 each contain all of $[x, y]$ and so $[x, y] \subseteq \mathcal{F}_1 \cap \mathcal{F}_2$. We therefore see that $\mathcal{F}_1 \cap \mathcal{F}_2$ is indeed a face of $\mathcal{C}_1 \cap \mathcal{C}_2$. ■

Although each pseudo-convex set \mathcal{C} comes with an entire structure of associated faces and sub-faces, we observe from the final point of Lemma 5.1.5 that the larger faces of \mathcal{C} determine the structure of the smaller faces of \mathcal{C} contained within them. Thus we will focus our attention on maximal faces, as defined below.

Definitions 5.1.7 (Maximal Faces and Maximal Convexity) Let \mathcal{C} be a pseudo-convex set. A *maximal face* of \mathcal{C} is a face of \mathcal{C} that is not contained within any other face of \mathcal{C} . Similarly, a *maximally convex* subset of \mathcal{C} is a convex subset of \mathcal{C} that is not contained within any other convex subset of \mathcal{C} .

We work now towards establishing an equivalence between maximal faces and maximally convex subsets, as seen in Theorem 5.1.9.

Lemma 5.1.8 If \mathcal{C} is any set then every convex subset of \mathcal{C} is contained within some maximally convex subset of \mathcal{C} .

Proof Let \mathcal{S} be any convex subset of \mathcal{C} and let P be the set of all convex subsets \mathcal{T} for which $\mathcal{S} \subseteq \mathcal{T} \subseteq \mathcal{C}$, noting that P is non-empty since we at least have $\mathcal{S} \in P$. Then (P, \subseteq) , i.e., the set P under the relation of set inclusion, forms a partially ordered set (poset).

Consider any chain Q in this poset, i.e., any subset $Q \subseteq P$ for which set inclusion imposes a total order, and let $\mathcal{M} = \bigcup_{\mathcal{T} \in Q} \mathcal{T}$. We claim that $\mathcal{M} \in P$ and that \mathcal{M} is an upper bound for the chain Q .

Take any points $x, y \in \mathcal{M}$. Then $x \in \mathcal{T}_1$ and $y \in \mathcal{T}_2$ for some $\mathcal{T}_1, \mathcal{T}_2 \in Q$. Since Q is totally ordered we have either $\mathcal{T}_1 \subseteq \mathcal{T}_2$ or $\mathcal{T}_2 \subseteq \mathcal{T}_1$; without loss of generality say $\mathcal{T}_1 \subseteq \mathcal{T}_2$. So $x, y \in \mathcal{T}_2$ and by the convexity of \mathcal{T}_2 we have the entire interval $[x, y]$ contained in \mathcal{T}_2 . Thus $[x, y] \subseteq \mathcal{M}$ and we see that \mathcal{M} is convex. Furthermore since each $\mathcal{T} \in Q$ satisfies $\mathcal{S} \subseteq \mathcal{T} \subseteq \mathcal{C}$ we have $\mathcal{S} \subseteq \mathcal{M} \subseteq \mathcal{C}$, and so $\mathcal{M} \in P$. In addition it is clear that for any $\mathcal{T} \in Q$ we have $\mathcal{T} \subseteq \mathcal{M}$ and thus \mathcal{M} is an upper bound for the chain Q .

So (P, \subseteq) is a non-empty poset for which every chain has an upper bound. By Zorn's lemma it follows that P has some maximal element \mathcal{F} . In particular this means that \mathcal{F} is convex with $\mathcal{S} \subseteq \mathcal{F} \subseteq \mathcal{C}$ and that no other convex set \mathcal{F}' can satisfy $\mathcal{F} \subset \mathcal{F}' \subseteq \mathcal{C}$. So \mathcal{F} is a maximally convex subset of \mathcal{C} containing \mathcal{S} as required. ■

Theorem 5.1.9 If \mathcal{C} is a pseudo-convex set then the maximal faces of \mathcal{C} are precisely the maximally convex subsets of \mathcal{C} .

Proof Let \mathcal{F} be any maximal face of \mathcal{C} . Since \mathcal{F} is convex, Lemma 5.1.8 shows that $\mathcal{F} \subseteq \mathcal{M}$ for some maximally convex subset $\mathcal{M} \subseteq \mathcal{C}$. But from Lemma 5.1.5 we see that \mathcal{M} is a face of \mathcal{C} and so the maximality of face \mathcal{F} implies that $\mathcal{F} = \mathcal{M}$. Thus \mathcal{F} is maximally convex.

Conversely, let \mathcal{F} be any maximally convex subset of \mathcal{C} . Lemma 5.1.5 shows that \mathcal{F} is a face of \mathcal{C} , and since all faces are convex the maximal convexity of \mathcal{F} implies that it cannot be contained within any other face. Thus \mathcal{F} is a maximal face of \mathcal{C} . ■

Corollary 5.1.10 If \mathcal{P} is a convex set then the one and only maximal face of \mathcal{P} is \mathcal{P} itself.

Proof This is immediate from Theorem 5.1.9 since the only maximally convex subset of \mathcal{P} is \mathcal{P} itself. ■

The remarks preceding Definitions 5.1.7 suggest that the maximal faces of a pseudo-convex set \mathcal{C} might in fact be sufficient to determine the entire face structure of \mathcal{C} . We can now show this to be true through the following result.

Lemma 5.1.11 Each pseudo-convex set is the union of its maximal faces.

Proof Let \mathcal{C} be any pseudo-convex set. From Lemma 5.1.8 we see that each point of \mathcal{C} (which itself forms a convex set) belongs to some maximally convex set of \mathcal{C} which from Theorem 5.1.9 is in fact a maximal face of \mathcal{C} . Since \mathcal{C} is the union of all its points it is therefore true that \mathcal{C} is also the union of its maximal faces. ■

So it follows from Lemma 5.1.11 and the results of Lemma 5.1.5 (and in particular from the last result of Lemma 5.1.5) that an understanding of the maximal faces of any pseudo-convex set \mathcal{C} is sufficient to determine the entire structure of faces and sub-faces of \mathcal{C} .

In closing this section we return our attention to the projective solution space for normal surfaces, the analysis of which is the primary motivation behind this and the following chapter.

Definition 5.1.12 (Embedded Faces) A face of the projective solution space (which is a polytope as described earlier) is an *embedded face* if every point on the face (including its boundary) has at most one non-zero quadrilateral coordinate associated with any tetrahedron.

Recall that the quadrilateral condition described above, i.e., that a point have at most one non-zero quadrilateral coordinate associated with any tetrahedron, is precisely the additional condition required of a normal surface vector if it is to represent an embedded normal surface. Likewise it follows that this condition is the precise additional condition required of a rational vector in the projective solution space if it is to represent a family of embedded normal surfaces.

Thus in order to understand the embedded normal surfaces within a triangulation it suffices to understand the embedded faces of the projective solution space.

Lemma 5.1.13 Any sub-face of an embedded face is itself an embedded face.

Proof This is clear from Definition 5.1.12 since every point of a sub-face is also a point of the original face. ■

Lemma 5.1.14 Let \mathcal{F} be a face of the projective solution space and let \mathbf{e} be some point in the relative interior of \mathcal{F} for which \mathbf{e} has at most one non-zero quadrilateral coordinate associated with any tetrahedron. Then the entire face \mathcal{F} is an embedded face.

Proof Suppose this is not true. Then there is some point $\mathbf{n} \in \mathcal{F}$ that does not satisfy the above quadrilateral coordinate condition. In particular there must be some coordinate position q for which $e_q = 0$ and $n_q > 0$ (where e_q and n_q are the q th coordinates of \mathbf{e} and \mathbf{n} respectively). Furthermore, since faces are convex and \mathbf{e} is within the face interior, vector $\mathbf{v} = \mathbf{e} - \varepsilon(\mathbf{n} - \mathbf{e})$ must also belong to face \mathcal{F} for some small $\varepsilon > 0$.

If we examine the q th coordinate of \mathbf{v} however, we see that $v_q = e_q - \varepsilon(n_q - e_q) = -\varepsilon n_q < 0$, contradicting the known fact that all points of the projective solution space have entirely non-negative coordinates. ■

Definition 5.1.15 (Maximal Embedded Faces) A *maximal embedded face* of the projective solution space is an embedded face that is not a sub-face of any other embedded face.

Lemma 5.1.16 The set of all points in the projective solution space having at most one non-zero quadrilateral coordinate associated with any tetrahedron is in fact the union of all maximal embedded faces.

Proof From Lemma 5.1.14 we see that the set of all points satisfying the above quadrilateral condition is the union of all embedded faces of the projective solution space. Furthermore, since the projective solution space is a polytope formed from finitely many matching equations, this polytope has finitely many faces. So every embedded face must be a sub-face of some maximal embedded face and we have the desired result. ■

Thus we see that the maximal embedded faces tell us all we wish to know about embedded normal surfaces. It is therefore upon the number of maximal embedded faces that we seek to place an upper bound.

5.2 Edge Weight Space

Instead of using the standard triangle and quadrilateral coordinates described earlier to represent normal surface as vectors in \mathbb{R}^{7n} , we shift now to a different set of coordinates. These are *edge weight coordinates*, used to represent normal surfaces as vectors in the smaller space \mathbb{R}^{n+1} . Edge weight coordinates have been described by Andrew Casson in seminars and informal settings, although no development of them has been published to date.

The edge weight coordinates described below have the advantage that they can only represent embedded normal surfaces. Thus there is no need to sift through solutions to the edge weight matching equations to determine which surfaces are embedded, since this is true of all solutions. Furthermore the underlying vector space of \mathbb{R}^{n+1} is much smaller, resulting in a more streamlined representation of the solution space.

The primary disadvantage with edge weight coordinates is that the solution space is no longer convex as it is with standard coordinates (though it is seen later in this chapter to be pseudo-convex). This means that computing vertex normal surfaces using edge weight coordinates becomes difficult since the standard algorithms used with most coordinate systems (including standard coordinates in \mathbb{R}^{7n} and quadrilateral-only coordinates in \mathbb{R}^{3n}) all rely on this convexity.

Nevertheless, this loss of convexity and the small underlying vector space impose an interesting geometric structure on the edge weight solution space. This structure is used to our advantage in the complexity calculations that follow.

Lemma 5.2.1 Any triangulation of a closed 3-manifold with one vertex and n tetrahedra has precisely $n + 1$ edges.

Proof This is a simple Euler characteristic calculation. A closed triangulation with n tetrahedra has $2n$ faces (since the $4n$ tetrahedron faces are identified in pairs) and has Euler characteristic 0. Thus $1 - E + 2n - n = 0$ where E is the number of edges, and so $E = n + 1$. ■

Definition 5.2.2 (Edge Weight Representation) Let T be a triangulation with distinct edges e_1, e_2, \dots, e_{n+1} and let S be an embedded normal surface in this triangulation. Then the *edge*

weight representation of S is the non-negative integer vector $(x_1, x_2, \dots, x_{n+1}) \in \mathbb{R}^{n+1}$ where x_i is the number of points in which S intersects edge e_i .

Theorem 5.2.3 Let T be a triangulation with distinct edges e_1, e_2, \dots, e_{n+1} and let S be an embedded normal surface in this triangulation with edge weight representation $\mathbf{x} = (x_1, x_2, \dots, x_{n+1})$. Then the following properties are true of \mathbf{x} .

- **Matching Equations:** Consider any tetrahedron Δ of T . Say the six edges of Δ are e_a, e_A, e_b, e_B, e_c and e_C (note that these edges need not be distinct edges of T), and say that edges e_a, e_b and e_c are opposite edges e_A, e_B and e_C respectively in tetrahedron Δ . Then precisely one of the following statements is true.
 - $x_a + x_A = x_b + x_B = x_c + x_C$ and S has no quadrilateral discs within Δ ;
 - $x_b + x_B = x_c + x_C > x_a + x_A$ and S has one or more quadrilaterals within Δ separating edges e_a and e_A ;
 - $x_c + x_C = x_a + x_A > x_b + x_B$ and S has one or more quadrilaterals within Δ separating edges e_b and e_B ;
 - $x_a + x_A = x_b + x_B > x_c + x_C$ and S has one or more quadrilaterals within Δ separating edges e_c and e_C .
- **Face Inequalities:** Consider any face of T . Say the three edges bordering this face are e_i, e_j and e_k (again note that these edges need not be distinct edges of T). Then $x_i \leq x_j + x_k$, $x_j \leq x_k + x_i$ and $x_k \leq x_i + x_j$.
- **Face Parity Condition:** Consider any face of T . Say the three edges bordering this face are e_i, e_j and e_k . Then the integer $x_i + x_j + x_k$ is even.

Note that the phrase *edge weight matching equations* refers only to the first (and not all) of these properties.

Proof For the matching equations, it is clear that no more than one of the given statements can be true. Furthermore, note that each triangular disc in tetrahedron Δ touches precisely one of e_a and e_A , precisely one of e_b and e_B and precisely one of e_c and e_C . Thus if S has no quadrilaterals in Δ and t triangles in Δ , we see that $x_a + x_A = x_b + x_B = x_c + x_C = t$.

If on the other hand S contains one or more quadrilaterals in Δ , the fact that S is embedded implies that it can have quadrilaterals of only one type. Let this be the type of quadrilateral separating edges e_c and e_C , and let there be q such quadrilaterals.

Then we see that $x_a + x_A = x_b + x_B = t + 2q$ since each quadrilateral touches each of e_a, e_A, e_b and e_B precisely once, and that $x_c + x_C = t$ since each quadrilateral misses e_c and e_C entirely. Hence $x_a + x_A = x_b + x_B > x_c + x_C$.

The remaining two statements are similarly established as consequences of using either of the remaining two quadrilateral types.

For the face inequalities and the face parity condition, let surface S have a_1, a_2 and a_3 normal arcs of each type crossing some face bounded by edges e_i, e_j and e_k as illustrated in Figure 5.3. It is clear from the diagram that $a_1 = \frac{1}{2}(x_j + x_k - x_i)$, $a_2 = \frac{1}{2}(x_k + x_i - x_j)$ and $a_3 = \frac{1}{2}(x_i + x_j - x_k)$.




Figure 5.3: Normal arcs crossing a tetrahedron face

The face inequalities and the face parity condition are then immediate consequences of the fact that a_1 , a_2 and a_3 are each non-negative integers. ■

As it happens, the properties described in Theorem 5.2.3 are sufficient to completely characterise the edge weight solution space.

Theorem 5.2.4 Let T be a triangulation with distinct edges e_1, e_2, \dots, e_{n+1} and let $\mathbf{x} \in \mathbb{R}^{n+1}$ be a vector of non-negative integers satisfying the constraints listed in Theorem 5.2.3. Then \mathbf{x} is the edge weight representation of a unique embedded normal surface in T .

Proof Consider some $\mathbf{x} \in \mathbb{R}^{n+1}$ satisfying the matching equations, the face inequalities and the face parity condition as described in Theorem 5.2.3, and assume we have some embedded normal surface S for which \mathbf{x} is the edge weight representation.




Figure 5.4: Normal discs in a tetrahedron

Consider any tetrahedron Δ of T with edges e_a, e_A, e_b, e_B, e_c and e_C as illustrated in Figure 5.4. Applying Theorem 5.2.3 to S we see that the type of the quadrilateral discs contained in S , if there are any at all, can be completely determined by the relative ordering of $x_a + x_A$, $x_b + x_B$ and $x_c + x_C$. Without loss of generality assume $x_b + x_B = x_c + x_C \geq x_a + x_A$, so S has either no quadrilateral discs in Δ or only discs of the type that separate edges e_a and e_A , again illustrated in Figure 5.4.

This figure labels the four types of triangular disc as 1, 2, 3 and 4. Let there be t_1, t_2, t_3 and t_4 triangular discs of each respective type and q quadrilateral discs in Δ , where $t_1, t_2, t_3, t_4, q \geq 0$. It is clear from the diagram that the following relations hold.

$$\begin{array}{lll} x_a = t_1 + t_4 & x_b = t_2 + t_4 + q & x_c = t_3 + t_4 + q \\ x_A = t_2 + t_3 & x_B = t_1 + t_3 + q & x_C = t_1 + t_2 + q \end{array}$$

From these we can create a variety of expressions for t_1, t_2, t_3 and t_4 in terms of the coordinates of \mathbf{x} .

$$\begin{aligned} t_1 &= \frac{1}{2}(x_a + x_B - x_c) & t_2 &= \frac{1}{2}(x_A + x_b - x_c) & t_3 &= \frac{1}{2}(x_A + x_B - x_C) & t_4 &= \frac{1}{2}(x_a + x_b - x_C) \\ t_1 &= \frac{1}{2}(x_a + x_C - x_b) & t_2 &= \frac{1}{2}(x_A + x_C - x_B) & t_3 &= \frac{1}{2}(x_A + x_c - x_b) & t_4 &= \frac{1}{2}(x_a + x_c - x_B) \end{aligned}$$

Furthermore we can create expressions for q and for the various sums $t_i + q$.

$$\begin{aligned} t_1 + q &= \frac{1}{2}(x_B + x_C - x_A) & t_2 + q &= \frac{1}{2}(x_b + x_C - x_a) \\ t_3 + q &= \frac{1}{2}(x_B + x_c - x_a) & t_4 + q &= \frac{1}{2}(x_b + x_c - x_A) \\ q &= \frac{1}{2}(x_b + x_B - x_a - x_A) & q &= \frac{1}{2}(x_c + x_C - x_a - x_A) \end{aligned}$$

From these equations it is clear that if S exists then it is unique, since the number of discs of each type is uniquely determined by the coordinates of \mathbf{x} .

Furthermore, since $x_b + x_B = x_c + x_C$ we see that these equations are not inconsistent and so whether or not such a surface S exists we can nevertheless construct a vector $\mathbf{s} \in \mathbb{R}^{7n}$ consisting of the (hypothetical) triangular and quadrilateral disc counts t_i and q for each tetrahedron satisfying all of the above expressions. If we can show that each coordinate of \mathbf{s} is a non-negative integer and that \mathbf{s} satisfies the standard normal surface matching equations then it will follow from Theorem 1.1.8 that such an embedded normal surface S does indeed exist.

It is immediate from the face parity condition that each of the values t_i and $t_i + q$ described above are integers, from which it follows that q is also an integer. From the face inequalities we see that each t_i is non-negative, and from the condition $x_b + x_B \geq x_a + x_A$ we see that q is non-negative.

Finally we attend to the standard normal surface matching equations. Recall from Theorem 1.1.7 that for each non-boundary face F of the underlying triangulation there are three such equations, one for each type of normal arc on F . Specifically such an equation states that the number of normal discs on one side of F meeting F in any particular type of normal arc must equal the number of normal discs on the other side of F meeting F in the same type of normal arc.

So consider face XYZ in Figure 5.4, and in particular consider the side of this face represented by the illustrated tetrahedron. The number of discs meeting face XYZ in a normal arc running past vertex X is $t_3 + q$, which using the above relations equates to $\frac{1}{2}(x_B + x_c - x_a)$. The number of discs meeting face XYZ in an arc running past vertex Y is just $t_4 = \frac{1}{2}(x_a + x_c - x_B)$, and the number of discs meeting face XYZ in an arc running past vertex Z is $t_1 = \frac{1}{2}(x_a + x_B - x_c)$.

What we find then is that for any face F bounded by edges e_i, e_j and e_k and any side of this face, the number of discs meeting F in a normal arc opposite edge e_k is simply $\frac{1}{2}(x_i + x_j - x_k)$ (and similarly for the other two types of normal arc). In particular this number of discs is independent of which side of F we are examining, and so the standard normal surface matching equations are satisfied by \mathbf{x} . Thus as explained above a unique embedded normal surface S exists for which \mathbf{x} is the edge weight representation. ■

As with standard triangle and quadrilateral coordinates, we find that whenever vector \mathbf{x} is an

edge weight representation for an embedded normal surface, vectors $k\mathbf{x}$ also represent (possibly disconnected) embedded normal surfaces for all positive integers k . To remove this redundancy in the solution space, we employ the same tool that is used with standard coordinates and define a corresponding projective solution space.

Definition 5.2.5 (Projective Edge Weight Solution Space) Let T be a triangulation with distinct edges e_1, e_2, \dots, e_{n+1} . The *projective edge weight solution space* of T is the set of all vectors $\mathbf{x} = (x_1, x_2, \dots, x_{n+1})$ in \mathbb{R}^{n+1} satisfying the following properties.

- The coordinates of \mathbf{x} are all non-negative reals (though they need not be integers);
- \mathbf{x} satisfies the edge weight matching equations and face inequalities as described in Theorem 5.2.3;
- $x_1 + x_2 + \dots + x_{n+1} = 1$.

The bounded region of the hyperplane $x_1 + x_2 + \dots + x_{n+1} = 1$ in \mathbb{R}^{n+1} consisting only of points with entirely non-negative coordinates that satisfy the face inequalities is called the *projection region* of \mathbb{R}^{n+1} and will be denoted by \mathbb{P}^n (since it will be seen to form a convex n -ball).

If \mathbf{x} is the edge weight representation of some embedded normal surface and $\lambda\mathbf{x}$ lies in the projective edge weight solution space for some $\lambda \in \mathbb{R}$ then $\lambda\mathbf{x}$ is called a *projective edge weight representation* of this surface (or of \mathbf{x}).

In cases where confusion might arise, the original projective solution space constructed using triangle and quadrilateral coordinates in \mathbb{R}^{7n} will be referred to as the *projective standard solution space*.

Lemma 5.2.6 The projection region \mathbb{P}^n is a finite closed convex n -dimensional polytope. Furthermore the vector $\mathbf{o} = (\frac{1}{n+1}, \frac{1}{n+1}, \dots, \frac{1}{n+1}) \in \mathbb{R}^{n+1}$ lies in the relative interior of \mathbb{P}^n .

Proof Let $\mathbf{x} = (x_1, x_2, \dots, x_{n+1}) \in \mathbb{P}^n$. Then $x_i \geq 0$ for each i , but since $\sum x_i = 1$ we also have $x_i \leq 1$ for each i and so $\mathbf{x} \in [0, 1]^{n+1}$. Thus \mathbb{P}^n is finite.

Let \mathcal{H} denote the hyperplane $x_1 + x_2 + \dots + x_{n+1} = 1$ in \mathbb{R}^{n+1} . Then each face inequality and each inequality $x_i \geq 0$ represents a closed convex half-space in \mathbb{R}^{n+1} for which \mathbb{P}^n is the intersection of \mathcal{H} with each of these finitely many half-spaces. Thus the finite space \mathbb{P}^n is in fact a closed convex polytope.

It is clear that $\mathbf{o} \in \mathcal{H}$. Let $\mathbf{h} = (h_1, h_2, \dots, h_{n+1})$ be any point in hyperplane \mathcal{H} ; we will examine the point $\mathbf{o} + \varepsilon(\mathbf{h} - \mathbf{o})$ for small $\varepsilon > 0$.

Note first that $\mathbf{o} + \varepsilon(\mathbf{h} - \mathbf{o})$ lies in \mathcal{H} since \mathcal{H} is a hyperplane. Furthermore, the i th coordinate of this vector is $\frac{1}{n+1} + \varepsilon(h_i - \frac{1}{n+1})$ which is strictly positive for small enough $\varepsilon > 0$. Finally each face inequality $x_i \leq x_j + x_k$ becomes $\frac{1}{n+1} + \varepsilon(h_i - \frac{1}{n+1}) \leq \frac{2}{n+1} + \varepsilon(h_j - \frac{1}{n+1}) + \varepsilon(h_k - \frac{1}{n+1})$ which is trivially satisfied for small enough $\varepsilon > 0$. So, since this list of constraints is finite, we find that $\mathbf{o} + \varepsilon(\mathbf{h} - \mathbf{o}) \in \mathbb{P}^n$ for small enough $\varepsilon > 0$.

Thus some small neighbourhood of \mathbf{o} in \mathcal{H} is contained entirely within \mathbb{P}^n . Since $\mathbb{P}^n \subseteq \mathcal{H}$ it follows that the dimension of \mathbb{P}^n is precisely the dimension of \mathcal{H} which is n , and that \mathbf{o} lies in the relative interior of \mathbb{P}^n . ■

We see then that, since the projective edge weight solution space is merely the set of all points in the finite polytope \mathbb{P}^n satisfying the edge weight matching equations, this reduces the matching equation solution space to a finite structure embedded in n dimensions.

Lemma 5.2.7 Let T be a triangulation. Then every non-empty embedded normal surface in T has a unique projective edge weight representation.

Furthermore, if \mathbf{x} is a vector in the projective edge weight solution space of T whose coordinates are all rational, then \mathbf{x} is the projective edge weight representation of some embedded normal surface in T .

Proof Let S be a non-empty embedded normal surface in T with edge weight representation \mathbf{x} . Then, since the matching equations and face inequalities are linear, Theorem 5.2.3 shows that $\lambda\mathbf{x}$ satisfies these equations and inequalities for all $\lambda > 0$. Furthermore, all coordinates of $\lambda\mathbf{x}$ are non-negative and sum to $\lambda(x_1 + x_2 + \dots + x_{n+1})$. Thus $\lambda = 1/(x_1 + x_2 + \dots + x_{n+1})$ is the unique $\lambda > 0$ for which $\lambda\mathbf{x}$ lies in the projective edge weight solution space (observing that the denominator of λ is non-zero because our surface is non-empty). Note that no $\lambda'\mathbf{x}$ lies in the projective edge weight solution space for $\lambda' \leq 0$ since the coordinate sum of $\lambda'\mathbf{x}$ would be non-positive (and so not 1).

Now let \mathbf{x} be some vector in the projective edge weight solution space with entirely rational coordinates. Let $k = 2m$ where m is the least common multiple of the denominators of the various coordinates of \mathbf{x} . Then $k\mathbf{x}$ has entirely even integer coordinates, and so the second result follows from Theorem 5.2.4. ■

From these results we see that the projective edge weight solution space allows a simple reconstruction of the entire set of embedded normal surfaces in a triangulation. We finish with an illustration before probing into the geometric structure of this solution space.

Example 5.2.8 Consider the triangulation of the lens space $L(5, 1)$ illustrated in Figure 5.5. Faces ACD and DBC are identified, faces WXY and XZW are identified, faces ABC and ZWY are identified and faces ABD and XYZ are identified. The three edges of this triangulation are labelled on the diagram as e_1 , e_2 and e_3 .




Figure 5.5: A triangulation of $L(5, 1)$

In the first tetrahedron then, the matching equations give us the following four options.

$$\begin{aligned} x_2 + x_2 &= x_3 + x_3 = x_3 + x_1; \\ x_2 + x_2 &= x_3 + x_3 > x_3 + x_1; \\ x_3 + x_3 &= x_3 + x_1 > x_2 + x_2; \\ x_3 + x_1 &= x_2 + x_2 > x_3 + x_3, \end{aligned}$$

which reduce to

$$\begin{aligned}x_1 &= x_2 = x_3; \\x_2 &= x_3 > x_1; \\x_3 &= x_1 > x_2; \\x_2 &= \frac{1}{2}(x_1 + x_3) \text{ and } x_1 > x_3.\end{aligned}$$

In the second tetrahedron the matching equations give us the following four options.

$$\begin{aligned}x_1 + x_1 &= x_3 + x_3 = x_3 + x_2; \\x_1 + x_1 &= x_3 + x_3 > x_3 + x_2; \\x_3 + x_3 &= x_3 + x_2 > x_1 + x_1; \\x_3 + x_2 &= x_1 + x_1 > x_3 + x_3,\end{aligned}$$

which reduce to

$$\begin{aligned}x_1 &= x_2 = x_3; \\x_3 &= x_1 > x_2; \\x_2 &= x_3 > x_1; \\x_1 &= \frac{1}{2}(x_2 + x_3) \text{ and } x_2 > x_3.\end{aligned}$$

To find the projective edge weight solution space, we begin by seeking all points in \mathbb{R}^3 that satisfy at least one equation from the first set, at least one equation from the second set and the additional equation $x_1 + x_2 + x_3 = 1$ (the face inequalities are dealt with afterwards). We take cases according to the relative ordering of x_1 and x_3 .

- Say $x_1 < x_3$. Then the only applicable equation from the first set is $x_2 = x_3 > x_1$. This equation is also found in the second set, and so we have the family of solutions

$$\left(\frac{1}{3} - 2t, \frac{1}{3} + t, \frac{1}{3} + t\right) \text{ for } 0 < t \leq \frac{1}{6}$$

where $0 < t$ to ensure $x_1 < x_3$ and $t \leq \frac{1}{6}$ to ensure $x_1 \geq 0$.

- Say $x_1 = x_3$. If $x_2 = x_3$ also then we have the single solution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

If $x_2 > x_3$ then the only way we can satisfy the second set of equations is to have $x_1 = \frac{1}{2}(x_2 + x_3)$. But then $x_2 > x_1 > x_3$ and we cannot satisfy the first set of equations.

If $x_2 < x_3$ then the second set must be satisfied by $x_3 = x_1 > x_2$. This equation also belongs to the first set and so we have the family of solutions

$$\left(\frac{1}{3} + t, \frac{1}{3} - 2t, \frac{1}{3} + t\right) \text{ for } 0 < t \leq \frac{1}{6}$$

where $0 < t$ to ensure $x_2 < x_3$ and $t \leq \frac{1}{6}$ to ensure $x_2 \geq 0$.

- Say $x_1 > x_3$. Then we must satisfy the first set of equations by having $x_2 = \frac{1}{2}(x_1 + x_3)$. But then $x_1 > x_2 > x_3$ and the second set cannot be satisfied.

With all options exhausted, we now turn to the face inequalities. From the four faces of our

triangulation the following face inequalities are obtained.

$$\begin{array}{lll}
x_1 \leq x_2 + x_3 & x_2 \leq 2x_3 & x_1 \leq 2x_3 \\
x_2 \leq x_3 + x_1 & x_3 \leq x_2 + x_3 & x_3 \leq x_1 + x_3 \\
x_3 \leq x_1 + x_2 & &
\end{array}$$

Each of these inequalities is satisfied by all of the solutions described above. So we see that the projective solution space consists of two line segments (represented by the two parameterised families listed above) joined at the point $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. This projective solution space is illustrated in Figure 5.6, where the enclosing hyperplane $x_1 + x_2 + x_3 = 1$ is also shaded.




Figure 5.6: The projective edge weight solution space for the two-tetrahedron $L(5, 1)$

Note that the central point $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is in fact a multiple of $(2, 2, 2)$ which is the edge weight representation of the vertex linking sphere.

5.3 Geometric Structures

Recall the ultimate aim of this chapter, which is to bound the number of maximal embedded faces of the projective solution space. Our approach to this problem is to convert it into a problem regarding the projective edge weight solution space. With this in mind, we examine the various geometric structures that are formed by the edge weight matching equations in \mathbb{R}^{n+1} .

Definitions 5.3.1 An n -fan is the n -complex formed by joining three half-spaces of \mathbb{R}^n along their common \mathbb{R}^{n-1} boundary, as illustrated in Figure 5.7. Each individual half-space is called a *leaf* of the fan and the common \mathbb{R}^{n-1} boundary is called its *axis*.

Each leaf is defined to include its boundary (i.e., the axis). A *strict leaf* refers to a leaf with this axis excluded.

Definitions 5.3.2 Consider an n -fan linearly embedded in \mathbb{R}^{n+1} , so that each leaf forms half of an n -dimensional hyperplane and the axis forms an entire $(n - 1)$ -dimensional hyperplane. This n -fan divides the surrounding \mathbb{R}^{n+1} into three regions.

If each of these regions is convex and if no two leaves are parallel (i.e., aligned so that together they form an entire n -dimensional hyperplane) then this n -fan is referred to as *balanced*. Otherwise this n -fan is referred to as *unbalanced*.




Figure 5.7: A sample 2-fan

Example 5.3.3 The leftmost fan in Figure 5.8 is balanced; the centre fan and the rightmost fan are both unbalanced.




Figure 5.8: Balanced and unbalanced 2-fans in \mathbb{R}^3

Another way of envisaging balanced and unbalanced fans is to consider a 2-dimensional plane perpendicular to the axis of the fan. The axis meets this plane in a single point and the leaves meet this plane in rays emanating from this point, as illustrated in Figure 5.9. In a balanced fan these three rays divide the plane into regions whose angles are all $< 180^\circ$ (as illustrated), whereas an unbalanced fan has some region whose angle is $\geq 180^\circ$.




Figure 5.9: A planar cross-section of a balanced fan

Lemma 5.3.4 Consider an n -fan \mathcal{F} linearly embedded in \mathbb{R}^{n+1} . Then the following statements are equivalent, where all hyperplanes referred to are n -dimensional.

1. \mathcal{F} is balanced;
2. For each leaf \mathcal{L} of \mathcal{F} , the remaining two leaves of \mathcal{F} lie on opposite sides of the hyperplane containing \mathcal{L} ;
3. For any hyperplane \mathcal{H} containing the axis of \mathcal{F} , there is a leaf of \mathcal{F} on either side of \mathcal{H} .

Proof We prove this in three stages.

- $1 \Rightarrow 3$: We prove this in the contrapositive, i.e., $\neg 3 \Rightarrow \neg 1$. Suppose we have some fan \mathcal{F} and hyperplane \mathcal{H} for which the axis of \mathcal{F} lies within \mathcal{H} and all three leaves of \mathcal{F} lie within or on one fixed side of \mathcal{H} .

If two of these leaves lie within \mathcal{H} then these leaves are parallel, and so \mathcal{F} is unbalanced. Otherwise the region bounded by \mathcal{F} that includes the other side of \mathcal{H} , i.e., the side not containing any leaves, is non-convex and so again \mathcal{F} is unbalanced.

- $3 \Rightarrow 2$: This is immediate since statement 2 is merely a restriction of statement 3 requiring the hyperplane \mathcal{H} to be a hyperplane containing a leaf of \mathcal{F} .
- $2 \Rightarrow 1$: We again prove the contrapositive statement $\neg 1 \Rightarrow \neg 2$. Say fan \mathcal{F} is unbalanced. If two leaves \mathcal{L}_1 and \mathcal{L}_2 of \mathcal{F} are parallel then the hyperplane containing \mathcal{L}_1 also contains \mathcal{L}_2 , since both leaves meet along the axis of \mathcal{F} . This hyperplane therefore does not have a leaf of \mathcal{F} on either side.

Otherwise there is a non-convex region of \mathbb{R}^{n+1} bounded by two leaves \mathcal{L}_1 and \mathcal{L}_2 . Consider the hyperplane \mathcal{H} containing \mathcal{L}_1 ; this non-convex region includes the entire half of \mathbb{R}^{n+1} on one side of \mathcal{H} , and so the remaining two leaves of \mathcal{F} must both lie on the other side of \mathcal{H} .

Thus we see that all three statements are equivalent. ■

Lemma 5.3.5 A balanced n -fan \mathcal{F} in \mathbb{R}^{n+1} is pseudo-convex. Furthermore each non-strict leaf of \mathcal{F} forms a maximal face of \mathcal{F} , and the axis of \mathcal{F} also forms a face of \mathcal{F}

Proof Consider any two points $x, y \in \mathcal{F}$. We examine the possible placements of the open interval (x, y) .

- If x and y both lie on the axis of \mathcal{F} then (x, y) lies entirely within this axis.
- If x lies on a strict leaf of \mathcal{F} and y lies on this same strict leaf or on the axis of \mathcal{F} then (x, y) lies entirely within the strict leaf containing x .
- If x and y lie on different strict leaves of \mathcal{F} then let \mathcal{R} be the region of \mathbb{R}^{n+1} bounded by these two leaves. Since \mathcal{F} is balanced these leaves are non-parallel and \mathcal{R} is convex, and therefore the open interval (x, y) lies in the interior of region \mathcal{R} . In particular, (x, y) contains no points at all from \mathcal{F} .

Recall then the conditions for pseudo-convexity as presented in Definition 5.1.1. Consider some convex subset $\mathcal{S} \subseteq \mathcal{F}$ and two distinct points $x, y \in \mathcal{F}$ for which the open interval (x, y) meets \mathcal{S} . Note that since \mathcal{S} is convex and the leaves of \mathcal{F} are non-parallel, \mathcal{S} must lie entirely within some non-strict leaf \mathcal{L} . From the above discussion we have two cases to consider.

- Points x and y both lie on the axis of \mathcal{F} . In this case the entire closed interval $[x, y]$ lies within this axis and hence within \mathcal{L} .
- The interval (x, y) lies within some strict leaf of \mathcal{F} . In this case, since (x, y) meets $\mathcal{S} \subseteq \mathcal{L}$, this strict leaf must represent the same leaf as \mathcal{L} . Thus we see again that the closed interval $[x, y]$ lies entirely within the closure of this strict leaf, i.e., within non-strict leaf \mathcal{L} .

In each case we see that \mathcal{L} is a convex subset of \mathcal{F} containing both \mathcal{S} and $[x, y]$. Thus balanced fan \mathcal{F} is pseudo-convex.

Consider now any non-strict leaf \mathcal{L} of \mathcal{F} . Since fan \mathcal{F} is balanced, its leaves are non-parallel and so \mathcal{L} is a maximally convex subset of \mathcal{F} . From Theorem 5.1.9 it follows that \mathcal{L} is a maximal face of \mathcal{F} . Furthermore, since the axis of \mathcal{F} is the intersection of its three non-strict leaves it follows from Lemma 5.1.5 that this axis is also a face of \mathcal{F} . ■

Theorem 5.3.6 Consider a single tetrahedron Δ of an n -tetrahedron triangulation, and let \mathbf{o} be the special vector $(\frac{1}{n+1}, \frac{1}{n+1}, \dots, \frac{1}{n+1}) \in \mathbb{R}^{n+1}$.

The set of vectors in \mathbb{R}^{n+1} that satisfy the edge weight matching equations for only tetrahedron Δ forms one of the following structures in \mathbb{R}^{n+1} .

- A balanced n -fan whose axis passes through \mathbf{o} ;
- An n -dimensional hyperplane passing through \mathbf{o} ;
- A closed $(n + 1)$ -dimensional half-space of \mathbb{R}^{n+1} whose boundary is a hyperplane passing through \mathbf{o} ;
- The entire space \mathbb{R}^{n+1} .

Proof Let the six edges of Δ be e_a, e_A, e_b, e_B, e_c and e_C as described in Theorem 5.2.3, where $a, b, c, A, B, C \in \{1, 2, \dots, n + 1\}$ and edges e_a, e_b and e_c are opposite edges e_A, e_B and e_C respectively in Δ . Note that indices a, b, c, A, B and C need not be distinct. Let $\mathbf{x} = (x_1, x_2, \dots, x_{n+1})$ be an arbitrary vector in \mathbb{R}^{n+1} .

Define sums $\sigma_1 = x_a + x_A, \sigma_2 = x_b + x_B$ and $\sigma_3 = x_c + x_C$. Then the edge weight equations for Δ require one of the following statements to be true.

$$\begin{aligned} \sigma_1 &= \sigma_2 = \sigma_3; \\ \sigma_2 &= \sigma_3 > \sigma_1; \\ \sigma_3 &= \sigma_1 > \sigma_2; \\ \sigma_1 &= \sigma_2 > \sigma_3. \end{aligned}$$

We will take cases according to the relationships between sums σ_1, σ_2 and σ_3 . In the following arguments, all hyperplanes are n -dimensional.

- Consider first the case in which edge identifications within the triangulation cause σ_1, σ_2 and σ_3 to represent identical pairs of edges (for instance, if $a = b = c$ and $A = B = C$). Then $\sigma_1 = \sigma_2 = \sigma_3$ for all points in \mathbb{R}^{n+1} and so the vectors satisfying the matching equations for Δ form the entire space \mathbb{R}^{n+1} .
- Now say that two of these sums (but not the third) represent identical pairs of edges. Let these identical sums be σ_1 and σ_2 , so that $\sigma_1 = \sigma_2$ for all points in \mathbb{R}^{n+1} . Then our edge weight equations for Δ are satisfied if and only if $\sigma_1 = \sigma_3$ or $\sigma_1 > \sigma_3$. Since σ_1 and σ_3 are not identical sums, the equation $\sigma_1 = \sigma_3$ defines a hyperplane and so our solution space is the closed half-space $x_a + x_A \geq x_c + x_C$ which is bounded by this hyperplane $x_a + x_A = x_c + x_C$. Observe also that vector \mathbf{o} lies on this boundary hyperplane since $o_a = o_A = o_c = o_C = \frac{1}{n+1}$.

- Otherwise consider the case in which none of these sums are identical. Then equations $\sigma_1 = \sigma_2$, $\sigma_2 = \sigma_3$ and $\sigma_3 = \sigma_1$ each represent hyperplanes in \mathbb{R}^{n+1} , and vector \mathbf{o} can be seen to lie on each of these hyperplanes. Since all three hyperplanes intersect at \mathbf{o} , no two of these hyperplanes can be parallel unless they are in fact identical. Again we take cases.

- Suppose that two of these hyperplanes are identical; let these be $\sigma_1 = \sigma_2$ and $\sigma_2 = \sigma_3$. Then any point on both of these hyperplanes also satisfies $\sigma_3 = \sigma_1$ and we see that all three hyperplanes must be identical. Thus the only way a vector \mathbf{x} can satisfy the matching equations for Δ is if \mathbf{x} lies on this common hyperplane (and so $\sigma_1 = \sigma_2 = \sigma_3$). In this case the solution space is a single hyperplane which we have already observed passes through \mathbf{o} .
- Otherwise all three of these hyperplanes are distinct. Since these hyperplanes are pairwise non-parallel, it follows that any pair must intersect in an $(n - 1)$ -dimensional subspace of \mathbb{R}^{n+1} . However, any point lying on two of these hyperplanes must satisfy $\sigma_1 = \sigma_2 = \sigma_3$ and thus also lie on the third. So all three hyperplanes (as well as any isolated pair) intersect in the $(n - 1)$ -dimensional subspace defined by $\sigma_1 = \sigma_2 = \sigma_3$; call this common intersection \mathcal{A} .

The points in our solution set that satisfy $\sigma_2 = \sigma_3 > \sigma_1$ must be a half of the hyperplane $\sigma_2 = \sigma_3$ bounded by its intersection with the non-parallel hyperplane $\sigma_3 = \sigma_1$, i.e., an n -dimensional half-hyperplane bounded by the $(n - 1)$ -dimensional \mathcal{A} . Call this half-hyperplane \mathcal{L}_1 . Similarly, the points that satisfy $\sigma_3 = \sigma_1 > \sigma_2$ lie on another half-hyperplane bounded by \mathcal{A} (call this \mathcal{L}_2) and the points that satisfy $\sigma_1 = \sigma_2 > \sigma_3$ lie on yet another half-hyperplane bounded by \mathcal{A} (call this \mathcal{L}_3).

We see then that our entire solution set is an n -fan with leaves \mathcal{L}_1 , \mathcal{L}_2 and \mathcal{L}_3 and axis \mathcal{A} . We have already observed that \mathbf{o} lies on all three hyperplanes and so lies on axis \mathcal{A} . The only remaining task is to show that this n -fan is balanced.

Consider the hyperplane containing leaf \mathcal{L}_1 ; this is the hyperplane $f(\mathbf{x}) = 0$ where $f(\mathbf{x}) = \sigma_2 - \sigma_3$ (i.e., $f(\mathbf{x}) = x_b + x_B - x_c - x_C$). For any point \mathbf{x} lying on the strict leaf \mathcal{L}_2 we have $\sigma_3 > \sigma_2$ and so $f(\mathbf{x}) < 0$. For any point \mathbf{x} lying on the strict leaf \mathcal{L}_3 however we have $\sigma_2 > \sigma_3$ and so $f(\mathbf{x}) > 0$. Thus strict leaves \mathcal{L}_2 and \mathcal{L}_3 lie on opposite sides of the hyperplane containing \mathcal{L}_1 .

A similar argument can be used to show that any two strict leaves lie on opposite sides of the hyperplane containing the third, and so we see that our n -fan is indeed balanced.

Thus concludes our case analysis. ■

Since we plan to work in the projective edge weight solution space (as opposed to all of \mathbb{R}^{n+1}), it is important to understand how the geometric structures described in Theorem 5.3.6 intersect the projection region \mathbb{P}^n .

Lemma 5.3.7 Again let Δ be a single tetrahedron in an n -tetrahedron triangulation and let \mathcal{S} be the set of vectors in \mathbb{R}^{n+1} that satisfy the edge weight matching equations for only tetrahedron Δ . Let $\text{relint}(\mathbb{P}^n)$ denote the relative interior of the projection region \mathbb{P}^n .

If \mathcal{S} is a balanced n -fan then $\text{relint}(\mathbb{P}^n)$ contains points from all three strict leaves as well as from the axis of this fan. If \mathcal{S} is an $(n+1)$ -dimensional half-space with hyperplane boundary \mathcal{H} or \mathcal{S} is merely a single n -dimensional hyperplane \mathcal{H} then $\text{relint}(\mathbb{P}^n)$ contains points on either side of \mathcal{H} .

Proof As before let $\mathbf{o} = (\frac{1}{n+1}, \frac{1}{n+1}, \dots, \frac{1}{n+1}) \in \mathbb{R}^{n+1}$, and let \mathcal{H} be the hyperplane defined by $x_1 + x_2 + \dots + x_{n+1} = 1$ of which Lemma 5.2.6 tells us \mathbb{P}^n is a full-dimensional subspace. Recall also from Lemma 5.2.6 that $\mathbf{o} \in \text{relint}(\mathbb{P}^n)$ and observe that vector \mathbf{o} is in fact orthogonal to \mathcal{H} since the equation defining \mathcal{H} can be expressed as $\mathbf{o} \cdot (\mathbf{x} - \mathbf{o}) = 0$. Thus we can express every point $\mathbf{x} \in \mathbb{R}^{n+1}$ as $\mathbf{x} = \lambda \mathbf{o} + \mu(\mathbf{p} - \mathbf{o})$ for some $\mathbf{p} \in \text{relint}(\mathbb{P}^n)$ and some $\lambda, \mu \in \mathbb{R}$ where $\mu \geq 0$.

Suppose then that \mathcal{S} is a balanced n -fan. In this case we already know that $\text{relint}(\mathbb{P}^n)$ contains an axis point of this fan since $\mathbf{o} \in \text{relint}(\mathbb{P}^n)$. Let \mathcal{L} be one of the strict leaves of this fan and choose any $\mathbf{l} \in \mathcal{L}$.

As explained above we can express $\mathbf{l} = \lambda \mathbf{o} + \mu(\mathbf{p} - \mathbf{o})$ where $\mathbf{p} \in \text{relint}(\mathbb{P}^n)$ and $\mu \geq 0$. We recall from the proof of Theorem 5.3.6 that the equations defining the strict leaf \mathcal{L} are of the form

$$x_a + x_A = x_b + x_B > x_c + x_C \quad (5.1)$$

(possibly with the subscripts rearranged) and so in particular \mathbf{l} satisfies these equations. But since \mathbf{o} satisfies $x_a + x_A = x_b + x_B = x_c + x_C$ (since all of its coordinates are equal) it follows that $\mu \mathbf{p}$ satisfies Equations 5.1, and since $\mu > 0$ we see furthermore that \mathbf{p} satisfies these equations. Thus \mathbf{p} is a point of $\text{relint}(\mathbb{P}^n)$ belonging to strict leaf \mathcal{L} .

Consider now the case where \mathcal{S} is either an $(n+1)$ -dimensional half-space with hyperplane boundary \mathcal{H} or just a single n -dimensional hyperplane \mathcal{H} . Either way, we recall from the proof of Theorem 5.3.6 that the equation defining \mathcal{H} is of the form $x_a + x_A = x_b + x_B$, i.e.,

$$f(\mathbf{x}) = x_a + x_A - x_b - x_B = 0,$$

again possibly with different subscripts. Choose any point \mathbf{n} for which $f(\mathbf{n}) < 0$ and express $\mathbf{n} = \lambda \mathbf{o} + \mu(\mathbf{p} - \mathbf{o})$ for some $\mathbf{p} \in \text{relint}(\mathbb{P}^n)$ and $\mu \geq 0$. Once more \mathbf{o} lies on \mathcal{H} and so $f(\mathbf{o}) = 0$, giving $\mu f(\mathbf{p}) < 0$ and hence $f(\mathbf{p}) < 0$. Thus \mathbf{p} is a point of $\text{relint}(\mathbb{P}^n)$ on the negative side of \mathcal{H} . A point of $\text{relint}(\mathbb{P}^n)$ on the positive side of \mathcal{H} is similarly constructed. ■

Corollary 5.3.8 Consider a single tetrahedron Δ of an n -tetrahedron triangulation, and let \mathbf{o} be the special vector $(\frac{1}{n+1}, \frac{1}{n+1}, \dots, \frac{1}{n+1}) \in \mathbb{R}^{n+1}$. Then the set of vectors in the (convex n -dimensional) projection region \mathbb{P}^n that satisfy the edge weight matching equations for tetrahedron Δ forms one of the following structures in \mathbb{P}^n .

- A balanced $(n-1)$ -fan whose axis passes through \mathbf{o} ;
- An $(n-1)$ -dimensional hyperplane passing through \mathbf{o} ;
- A closed n -dimensional half-space of \mathbb{P}^n whose boundary hyperplane passes through \mathbf{o} ;
- The entire space \mathbb{P}^n .

Proof This follows immediately from Theorem 5.3.6 which describes the large-scale structure of the solution space for tetrahedron Δ in \mathbb{R}^{n+1} and Lemma 5.3.7 which ensures that this solution space does not lose its large scale structure when it is restricted to the projection region \mathbb{P}^n . ■

Thus we see that the entire projective edge weight solution space is merely an intersection of balanced $(n - 1)$ -fans, hyperplanes and half-spaces within \mathbb{P}^n , leading us to the following result.

Lemma 5.3.9 The projective edge weight solution space for any triangulation is finite, closed and pseudo-convex.

Proof Let the projective edge weight solution space be \mathcal{S} . Recall from Definition 5.2.5 that \mathcal{S} is the set of all points in the projection region \mathbb{P}^n that satisfy the edge weight matching equations for every tetrahedron in the triangulation, and so from Corollary 5.3.8 we see that \mathcal{S} is an intersection of finitely many balanced fans, hyperplanes and closed half-spaces in the convex polytope \mathbb{P}^n .

Lemma 5.2.6 tells us that \mathbb{P}^n is finite, closed and convex; furthermore every hyperplane and half-space is closed and convex, and every balanced fan is closed and pseudo-convex as seen in Lemma 5.3.5. Thus the intersection \mathcal{S} of each of these objects is finite, closed and (using Lemma 5.1.3) pseudo-convex. ■

5.4 Mappings Between Solution Spaces

Section 5.3 concentrates upon each geometric structure formed from the matching equations for a single tetrahedron. In Section 5.5 we proceed to combine these results and analyse in detail the geometric structure of the entire projective edge weight solution space.

In the meantime however, since we have established the pseudo-convexity of the projective edge weight solution space in Lemma 5.3.9, we can use Definition 5.1.4 to investigate its structure of faces and sub-faces. We thus pause to examine how the vectors and faces of this projective edge weight solution space correspond to the vectors and faces of the projective standard solution space described at the beginning of this chapter.

Definitions 5.4.1 (Edge Weight and Standard Maps) Consider some n -tetrahedron triangulation T with distinct edges e_1, e_2, \dots, e_{n+1} . Let S be the set of all vectors in \mathbb{R}^{7n} whose coordinates are all non-negative, which satisfy the standard normal surface matching equations for T as seen in Theorem 1.1.7, and for which at most one quadrilateral coordinate associated with each tetrahedron is non-zero. Let E be the set of all vectors in \mathbb{R}^{n+1} whose coordinates are all non-negative and which satisfy the edge weight matching equations and face inequalities for T as seen in Theorem 5.2.3. We define the following maps between S and E .

- The *edge weight map* for T is the map $\phi : S \rightarrow E$ defined as follows. Take any $\mathbf{v} \in S$. For any edge e_i of T , choose any tetrahedron Δ adjacent to e . Let t_i and t_2 represent the coordinates of \mathbf{v} corresponding to the two types of triangular disc within Δ that meet e_i , and let q_1 and q_2 represent the coordinates of \mathbf{v} corresponding to the two types of quadrilateral disc within Δ that meet e_i (note that since $\mathbf{v} \in S$ at most one of q_1 and q_2 can be non-zero). These four types of disc are illustrated in Figure 5.10. The coordinate x_i of $\phi(\mathbf{v})$ corresponding to edge e_i is then defined to be $x_i = t_1 + t_2 + q_1 + q_2$.




Figure 5.10: Defining the edge weight map




Figure 5.11: Defining the standard map

- The *standard map* for T is the map $\psi : E \rightarrow S$ defined as follows. Take any $\mathbf{x} \in E$ where $\mathbf{x} = (x_1, x_2, \dots, x_{n+1})$. For any tetrahedron Δ of T , let the edges of Δ be e_a, e_A, e_b, e_B, e_c and e_C arranged as illustrated in Figure 5.11. Since \mathbf{x} satisfies the edge weight matching equations, assume without loss of generality that $x_b + x_B = x_c + x_C \geq x_a + x_A$.

For the two quadrilateral types within Δ not shown in Figure 5.11, we define the corresponding coordinate of $\psi(\mathbf{x})$ to be 0. Let t_1, t_2, t_3 and t_4 be the coordinates of $\psi(\mathbf{x})$ corresponding to the triangular discs of types 1, 2, 3 and 4 as illustrated, and let q be the coordinate of $\psi(\mathbf{x})$ corresponding to the quadrilateral disc type illustrated (i.e., the disc type separating edges e_a and e_A). These coordinates of $\psi(\mathbf{x})$ are defined as follows.

$$\left. \begin{aligned} t_1 &= \frac{1}{2}(x_a + x_B - x_c) & t_1 &= \frac{1}{2}(x_a + x_C - x_b) \\ t_2 &= \frac{1}{2}(x_A + x_b - x_c) & t_2 &= \frac{1}{2}(x_A + x_C - x_B) \\ t_3 &= \frac{1}{2}(x_A + x_c - x_b) & t_3 &= \frac{1}{2}(x_A + x_B - x_C) \\ t_4 &= \frac{1}{2}(x_a + x_c - x_B) & t_4 &= \frac{1}{2}(x_a + x_b - x_C) \\ q &= \frac{1}{2}(x_b + x_B - x_a - x_A) & q &= \frac{1}{2}(x_c + x_C - x_a - x_A) \end{aligned} \right\} \quad (5.2)$$

Note that any of the above equations may be used; we will prove shortly that $\psi(\mathbf{x})$ is well-defined nevertheless.

It is straightforward to see from Definitions 5.4.1 that if N is a normal surface in T then ϕ and ψ convert its standard vector representation in \mathbb{R}^{7n} into its edge weight representation in \mathbb{R}^{n+1}

and vice versa. In this section however we are interested in the properties of ϕ and ψ as maps between the entire spaces S and E , most of whose vectors are not immediate representations of normal surfaces (since most of these vectors do not have integer coordinates).

Lemma 5.4.2 The edge weight map ϕ as described in Definitions 5.4.1 is well-defined, and for any $\mathbf{v} \in S$ it is true that $\phi(\mathbf{v}) \in E$ as claimed.

Proof The only choices to be made when evaluating $\phi(\mathbf{v})$ are the choices of which adjacent tetrahedron to examine for each edge. So consider some edge e_i with adjacent tetrahedra $\Delta_1, \dots, \Delta_d$. We can label these tetrahedra so that each Δ_j is adjacent to Δ_{j+1} along a face containing e_i , as illustrated in Figure 5.12. Let x_i be the coordinate of $\phi(\mathbf{v})$ corresponding to edge e_i .




Figure 5.12: Tetrahedra adjacent to edge e_i

Recall that if we choose Δ_j as the tetrahedron with which to calculate x_i , we obtain x_i as the sum of the four coordinates of \mathbf{v} corresponding to disc types in Δ_j adjacent to edge e_i . However, since \mathbf{v} satisfies the standard normal surface matching equations, we can apply these matching equations to the face joining Δ_j and Δ_{j+1} to see that x_i is also the sum of the four coordinates of \mathbf{v} corresponding to disc types in Δ_{j+1} adjacent to edge e_i . Applying this argument throughout the list of tetrahedra $\Delta_1, \dots, \Delta_d$ we see that x_i is independent of the choice of tetrahedron used and so ϕ is well-defined.

Furthermore it is clear that each coordinate of $\phi(\mathbf{v})$ is non-negative. Consider then any tetrahedron Δ with edges e_a, e_A, e_b, e_B, e_c and e_C as illustrated in Figure 5.11. Since \mathbf{v} can have at most one non-zero quadrilateral coordinate associated with Δ , let this if it exists be the coordinate corresponding to the quadrilateral type separating edges e_a and e_A , again as illustrated in Figure 5.11. Let $q \geq 0$ be this quadrilateral coordinate and let t_1, t_2, t_3 and t_4 be the coordinates of \mathbf{v} corresponding to the triangular discs of types 1, 2, 3 and 4 as shown in the figure. We see then that the coordinates of $\phi(\mathbf{v})$ corresponding to the illustrated edges are as follows.

$$\left. \begin{array}{lll} x_a = t_1 + t_4 & x_b = t_2 + t_4 + q & x_c = t_3 + t_4 + q \\ x_A = t_2 + t_3 & x_B = t_1 + t_3 + q & x_C = t_1 + t_2 + q \end{array} \right\} \quad (5.3)$$

In particular we see that $x_b + x_B = x_c + x_C \geq x_a + x_A$ and so the edge weight matching equations for Δ are satisfied. Furthermore consider the face containing edges e_a, e_c and e_B . We see then that

$$x_a + x_c - x_B = 2t_4 \quad x_c + x_B - x_a = 2(t_3 + q) \quad x_B + x_a - x_c = 2t_1$$

and so the face inequalities for this face are satisfied. A similar argument shows that the face inequalities are satisfied for each face of Δ . Applying these arguments to each tetrahedron Δ we finally conclude that $\phi(\mathbf{v}) \in E$. ■

Lemma 5.4.3 The standard map ψ as described in Definitions 5.4.1 is well-defined, and for any $\mathbf{x} \in E$ it is true that $\psi(\mathbf{x}) \in S$ as claimed.

Proof Let $\mathbf{x} \in E$. Consider any tetrahedron Δ as illustrated in Figure 5.11, where we assume without loss of generality that $x_b + x_B = x_c + x_C \geq x_a + x_A$. We see in Equations 5.2 that the equations on the left are equivalent to the equations on the right since $x_b + x_B = x_c + x_C$.

So if in fact $x_b + x_B = x_c + x_C > x_a + x_A$, this shows the coordinates of $\psi(\mathbf{x})$ corresponding to tetrahedron Δ to be well-defined. If on the other hand $x_a + x_A = x_b + x_B = x_c + x_C$, we have three possible diagrams to choose from corresponding to the three relations

$$\begin{aligned} x_b + x_B &= x_c + x_C \geq x_a + x_A; \\ x_c + x_C &= x_a + x_A \geq x_b + x_B; \\ x_a + x_A &= x_b + x_B \geq x_c + x_C. \end{aligned}$$

In all cases Equations 5.2 define each quadrilateral coordinate of $\psi(\mathbf{x})$ associated with Δ to be zero, and the possible expressions for coordinates t_1, t_2, t_3 and t_4 become

$$\begin{array}{lll} t_1 = \frac{1}{2}(x_a + x_B - x_c) & t_1 = \frac{1}{2}(x_a + x_C - x_b) & t_1 = \frac{1}{2}(x_B + x_C - x_A) \\ t_2 = \frac{1}{2}(x_A + x_b - x_c) & t_2 = \frac{1}{2}(x_A + x_C - x_B) & t_2 = \frac{1}{2}(x_b + x_C - x_A) \\ t_3 = \frac{1}{2}(x_A + x_c - x_b) & t_3 = \frac{1}{2}(x_A + x_B - x_C) & t_3 = \frac{1}{2}(x_c + x_B - x_A) \\ t_4 = \frac{1}{2}(x_a + x_c - x_B) & t_4 = \frac{1}{2}(x_a + x_b - x_C) & t_4 = \frac{1}{2}(x_c + x_b - x_A) \end{array}$$

which are again all equivalent since $x_a + x_A = x_b + x_B = x_c + x_C$. Thus $\psi(\mathbf{x})$ is well-defined.

Examining Equations 5.2, the face inequalities applied to \mathbf{x} show each t_i to be non-negative and since $x_b + x_B = x_c + x_C \geq x_a + x_A$ we see that q is also non-negative. So the coordinates of $\psi(\mathbf{x})$ are all non-negative, and it is clear from Definitions 5.4.1 that at most one quadrilateral coordinates of $\psi(\mathbf{x})$ associated with each tetrahedron is non-zero.

To show that $\psi(x)$ satisfies the standard normal surface matching equations, we employ an identical argument to that seen at the end of the proof of Theorem 5.2.4. More specifically we use an identical argument to show that for any side of any face F in the triangulation, the sum of the coordinates of $\psi(\mathbf{x})$ representing discs that meet F in any particular type of normal arc is a function only of the coordinates of \mathbf{x} corresponding to edges that bound F , and in particular is independent of which side of F we are examining. Thus we see that $\psi(\mathbf{x})$ satisfies the standard normal surface matching equations and so $\psi(\mathbf{x}) \in S$ as required. ■

Lemma 5.4.4 The edge weight map ϕ and the standard map ψ are inverses of each other. Furthermore, for any non-negative $\lambda \in \mathbb{R}$ and any $\mathbf{v} \in S$ and $\mathbf{x} \in E$ it is true that $\phi(\lambda\mathbf{v}) = \lambda\phi(\mathbf{v})$ and $\psi(\lambda\mathbf{x}) = \lambda\psi(\mathbf{x})$.

Proof Let $\mathbf{v} \in S$. Consider any tetrahedron Δ as illustrated in Figure 5.11, where without loss of generality the non-zero quadrilateral coordinate of \mathbf{v} in Δ (if it exists) represents the quadrilateral type separating edges e_a and e_A . Let $q \geq 0$ be this quadrilateral coordinate and as usual let t_1, t_2, t_3 and t_4 be the coordinates of \mathbf{v} corresponding to the four triangular discs of Δ . Then the coordinates of $\phi(\mathbf{v})$ corresponding to the edges of Δ are given by Equations 5.3 (seen in the proof of Lemma 5.4.2).

From these equations it can be seen that $x_b + x_B = x_c + x_C \geq x_a + x_A$ and so the coordinates of $\psi(\phi(\mathbf{v}))$ corresponding to tetrahedron Δ are determined by Equations 5.2. Simple arithmetic shows that Equations 5.3 followed by Equations 5.2 restore our original values of t_1, t_2, t_3, t_4 and q and so we see that $\psi(\phi(\mathbf{v})) = \mathbf{v}$.

Consider now any $\mathbf{x} \in E$ and again let Δ be some tetrahedron in our triangulation as illustrated in Figure 5.11. Without loss of generality let $x_b + x_B = x_c + x_C \geq x_a + x_A$, so the coordinates of $\psi(\mathbf{x})$ corresponding to tetrahedron Δ are given by Equations 5.2 where q represents the quadrilateral type in Δ separating edges e_a and e_A . It follows then that the coordinates of $\phi(\psi(\mathbf{x}))$ corresponding to the edges of Δ are given by Equations 5.3 and again through simple arithmetic we see that these two sets of equations together restore our original edge weight coordinates. So $\phi(\psi(\mathbf{x})) = \mathbf{x}$.

Finally consider any strictly positive $\lambda \in \mathbb{R}$ and any $\mathbf{v} \in S$ and $\mathbf{x} \in E$. Since the equations and inequalities that define S and E are all linear we see that $\lambda\mathbf{v} \in S$ and $\lambda\mathbf{x} \in E$. Furthermore in Definitions 5.4.1 the inequalities that determine which cases to consider are all linear and the equations that define ϕ and ψ in each of these cases are again linear. So $\phi(\lambda\mathbf{v}) = \lambda\phi(\mathbf{v})$ and $\psi(\lambda\mathbf{x}) = \lambda\psi(\mathbf{x})$. In conclusion it is clear from Definitions 5.4.1 that $\phi(\mathbf{0}) = \mathbf{0}$ and $\psi(\mathbf{0}) = \mathbf{0}$ where $\mathbf{0}$ is the zero vector and so these relations hold for all non-negative λ . ■

Lemma 5.4.5 The edge weight map ϕ is linear on any embedded face of the projective standard solution space, and the standard map ψ is linear on any face of the projective edge weight solution space.

Proof Let \mathcal{F} be some embedded face of the projective standard solution space and consider any tetrahedron Δ . If there are two vectors $\mathbf{u}, \mathbf{v} \in \mathcal{F}$ each with a non-zero quadrilateral coordinate associated with Δ and for which these non-zero coordinates correspond to different types of quadrilateral in Δ , then the midpoint $\frac{1}{2}(\mathbf{u} + \mathbf{v})$ must have non-zero coordinates corresponding to both of these quadrilateral types. But faces are convex, so $\frac{1}{2}(\mathbf{u} + \mathbf{v}) \in \mathcal{F}$ and we see from Definition 5.1.12 that \mathcal{F} cannot be an embedded face.

Therefore each tetrahedron Δ has at most one type of quadrilateral for which non-zero coordinates are seen in \mathcal{F} , and thus the same set of linear equations are used to define $\phi(\mathbf{v})$ for all $\mathbf{v} \in \mathcal{F}$. Hence ϕ is linear on \mathcal{F} .

Consider now any face \mathcal{F} of the projective edge weight solution space and again take any tetrahedron with edges e_a, e_b and e_c opposite edges e_A, e_B and e_C respectively. Suppose that there are two different vectors in \mathcal{F} for which two different equations in the set

$$\begin{aligned} x_b + x_B &= x_c + x_C > x_a + x_A \\ x_c + x_C &= x_a + x_A > x_b + x_B \\ x_a + x_A &= x_b + x_B > x_c + x_C \end{aligned}$$

hold. Without loss of generality let these two different equations be the first two; then the midpoint of these two vectors must satisfy $x_c + x_C > x_a + x_A$ and $x_c + x_C > x_b + x_B$. So this midpoint cannot satisfy the edge weight matching equations and convex face \mathcal{F} cannot be a face of the projective edge weight solution space.

Thus for each tetrahedron Δ there is some single equation in the set

$$\begin{aligned} x_b + x_B = x_c + x_C &\geq x_a + x_A \\ x_c + x_C = x_a + x_A &\geq x_b + x_B \\ x_a + x_A = x_b + x_B &\geq x_c + x_C \end{aligned}$$

satisfied by all points in \mathcal{F} . So the same linear equations are used to define $\psi(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{F}$ and we see that ψ is linear on \mathcal{F} . ■

Lemma 5.4.6 Let S' be the set of all vectors in the projective standard solution space for which at most one quadrilateral coordinate associated with each tetrahedron is non-zero. Let E' represent the projective edge weight solution space. Define maps $\phi' : S' \rightarrow E'$ and $\psi' : E' \rightarrow S'$ by

$$\phi'(\mathbf{v}) = \frac{\phi(\mathbf{v})}{\|\phi(\mathbf{v})\|} \quad \text{and} \quad \psi'(\mathbf{x}) = \frac{\psi(\mathbf{x})}{\|\psi(\mathbf{x})\|}.$$

Then ϕ' and ψ' are inverse maps.

Proof Observe first that S' and E' simply represent all vectors in S and E respectively whose coordinates sum to 1. For each $\mathbf{v} \in S'$ and $\mathbf{x} \in E'$ we see from Lemma 5.4.4 that $\phi(\mathbf{v}) \neq \mathbf{0}$ and $\psi(\mathbf{x}) \neq \mathbf{0}$ and so maps ϕ' and ψ' are well-defined. Furthermore it is clear that the coordinates of $\phi'(\mathbf{v})$ and $\psi'(\mathbf{x})$ each sum to 1 and so $\phi'(\mathbf{v}) \in E'$ and $\psi'(\mathbf{x}) \in S'$ as claimed.

From Lemma 5.4.4 we see that for any $\mathbf{v} \in S'$ we have $\phi'(\mathbf{v}) = \lambda\phi(\mathbf{v})$ for some $\lambda > 0$ and then $\psi'(\phi'(\mathbf{v})) = \mu\psi(\phi(\mathbf{v})) = \mu\mathbf{v}$ for some $\mu > 0$. However, since the coordinates of $\psi'(\phi'(\mathbf{v}))$ and the coordinates of \mathbf{v} each sum to 1 it follows that $\psi'(\phi'(\mathbf{v})) = \mathbf{v}$. A similar argument shows that $\phi'(\psi'(\mathbf{x})) = \mathbf{x}$ for any $\mathbf{x} \in E'$. ■

Theorem 5.4.7 There is a 1-to-1 correspondence between the maximal embedded faces of the projective standard solution space and the maximal faces of the projective edge weight solution space.

Proof Let \mathcal{F} be any maximal embedded face of the projective standard solution space. Since \mathcal{F} is convex, Lemma 5.4.5 shows that ϕ maps \mathcal{F} to a convex subset of E , and from Lemma 5.4.6 it follows that ϕ' maps \mathcal{F} to a convex subset of the projective edge weight solution space. From Lemma 5.1.8 and Theorem 5.1.9 it follows that ϕ' maps \mathcal{F} into some maximal face of the projective edge weight solution space.

Alternatively let \mathcal{G} be any maximal face of the projective edge weight solution space. Again \mathcal{G} is convex and so Lemmas 5.4.5 and 5.4.6 show that ψ' maps \mathcal{G} to a convex subset of the projective standard solution space in which every vector has at most one non-zero quadrilateral coordinate associated with each tetrahedron. But since the projective standard solution space is a finite convex polytope, there is some face \mathcal{G}' of this polytope for which $\mathcal{G} \subseteq \mathcal{G}'$ and $\text{relint}(\mathcal{G}')$ contains

some point in \mathcal{G} . From Lemma 5.1.14 it follows that \mathcal{G}' is an embedded face of this polytope and thus belongs to some maximal embedded face of this polytope. So likewise we see that ψ' maps \mathcal{G} into some maximal embedded face of the projective standard solution space.

However, since ϕ' and ψ' are inverse maps and no maximal (embedded) face can contain any other maximal (embedded) face, it follows that ϕ' and ψ' map maximal embedded faces *onto* maximal faces and vice versa. Thus maps ϕ' and ψ' induce a 1-to-1 correspondence between the maximal embedded faces of the projective standard solution space and the maximal faces of the projective edge weight solution space as required. \blacksquare

Thus, in order to achieve our stated goal of bounding the number of maximal embedded faces of the projective standard solution space, it suffices to bound the number of maximal faces of the projective edge weight solution space instead.

5.5 Constructing the Solution Space

Returning to the geometric analysis of the projective edge weight solution space, we recall that Section 5.3 makes some progress towards understanding the geometric structures formed from the matching equations for each single tetrahedron. We wish now to pool these results and understand the geometric structure of the entire projective edge weight solution space.

Theorem 5.5.1 Let T be a one-vertex n -tetrahedron triangulation. Then the projective edge weight solution space for T can be formed as follows.

- Begin with a convex d -dimensional polytope for some $d \leq n$ with some privileged interior point \mathbf{o} ;
- Optionally place $\leq d$ balanced $(d - 1)$ -fans in this polytope, all of whose axes pass through \mathbf{o} , and restrict our space to their intersection;
- Optionally restrict the resulting space to its intersection with a convex cone whose vertex is \mathbf{o} .

Clearly this procedure is not guaranteed to produce the projective solution space however it is performed. We merely claim that there is some particular way of following it (i.e., some particular choice for d , the positions of the fans and so on) that will produce our projective solution space.

Proof To begin, we observe that the projective edge weight solution space is simply the intersection of the matching equation solution spaces for each individual tetrahedron in our triangulation. From Corollary 5.3.8 we already know the structures of these individual spaces.

So again let $\mathbf{o} = (\frac{1}{n+1}, \frac{1}{n+1}, \dots, \frac{1}{n+1}) \in \mathbb{R}^{n+1}$ and recall that \mathbf{o} is in the relative interior of the convex n -dimensional polytope \mathbb{P}^n . We can begin by ignoring the tetrahedra whose matching equations are satisfied by all of \mathbb{P}^n (since this will not alter the final intersection), and so by simply letting $d = n$ the following fact is observed from Corollary 5.3.8.

Type I Representation: The projective edge weight solution space is the intersection of $\leq d$ objects within a convex d -dimensional polytope, where $d \leq n$, \mathbf{o} is an interior point of this polytope and each of these objects is one of the following.

- A balanced $(d - 1)$ -fan whose axis passes through \mathbf{o} ;
- A $(d - 1)$ -dimensional hyperplane passing through \mathbf{o} ;
- A closed d -dimensional half-space whose boundary hyperplane passes through \mathbf{o} .

Suppose then that we have a type I representation of the projective edge weight solution space as described above, where in particular $k \geq 1$ of the objects are hyperplanes (and thus $d \geq k \geq 1$). From this we produce a simpler type I representation based on a polytope of dimension $d - 1$.

So let \mathcal{H} be one of these k hyperplanes. We examine the intersection of \mathcal{H} with each of the remaining objects in our type I representation. Let there be m remaining objects (so $m \leq d - 1$) and label these objects $\mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_m$.

- Suppose that \mathcal{O}_i is a balanced $(d - 1)$ -fan. We have three possible cases.
 - If hyperplane \mathcal{H} contains a leaf of \mathcal{O}_i (and hence also contains the axis), the intersection $\mathcal{H} \cap \mathcal{O}_i$ becomes a $(d - 1)$ -dimensional half-space in \mathcal{H} whose boundary is the axis of \mathcal{O}_i (which we recall passes through \mathbf{o}).
 - If hyperplane \mathcal{H} contains the axis of \mathcal{O}_i but not any of its leaves, then consider any leaf \mathcal{L} of \mathcal{O}_i . We see that \mathcal{L} and \mathcal{H} cannot be parallel (since they intersect at \mathbf{o} but do not represent the same hyperplane), so they can intersect in at most a single $(d - 2)$ -dimensional hyperplane within \mathcal{H} . The axis of \mathcal{O}_i however is such a $(d - 2)$ -dimensional hyperplane, and so the intersection $\mathcal{H} \cap \mathcal{L}$ must be this axis alone.
Thus \mathcal{H} meets the entire fan \mathcal{O}_i in only its axis and so the intersection $\mathcal{H} \cap \mathcal{O}_i$ is a $(d - 2)$ -dimensional hyperplane in \mathcal{H} (which again contains \mathbf{o}).
 - Otherwise hyperplane \mathcal{H} does not contain the axis of \mathcal{O}_i . In this case \mathcal{H} cannot be parallel to any of the leaves of \mathcal{O}_i and so \mathcal{H} simply takes a cross-section of this fan. Thus we see that the intersection $\mathcal{H} \cap \mathcal{O}_i$ is a balanced $(d - 2)$ -fan in the $(d - 1)$ -dimensional space \mathcal{H} . Note also that since \mathbf{o} lies both on the axis of \mathcal{O}_i and on \mathcal{H} , it must likewise lie on the axis of the new fan $\mathcal{H} \cap \mathcal{O}_i$.

- Suppose that \mathcal{O}_i is a $(d - 1)$ -dimensional hyperplane passing through \mathbf{o} . If $\mathcal{O}_i = \mathcal{H}$ then the intersection $\mathcal{H} \cap \mathcal{O}_i$ is simply all of \mathcal{H} .

Otherwise \mathcal{O}_i cannot be parallel to \mathcal{H} (since they both contain \mathbf{o}) and so the intersection $\mathcal{H} \cap \mathcal{O}_i$ is a $(d - 2)$ -dimensional hyperplane within the $(d - 1)$ -dimensional space \mathcal{H} that passes through \mathbf{o} .

- Finally suppose that \mathcal{O}_i is a closed d -dimensional half-space whose boundary contains \mathbf{o} . Since \mathcal{H} also contains \mathbf{o} , it follows that if \mathcal{H} is parallel to this boundary then \mathcal{H} is in fact equal to this boundary and so the intersection $\mathcal{H} \cap \mathcal{O}_i$ is simply the hyperplane \mathcal{H} .

Otherwise the boundary hyperplane of \mathcal{O}_i is not parallel to \mathcal{H} and so it must intersect \mathcal{H} in a $(d - 2)$ -dimensional hyperplane \mathcal{I} within the $(d - 1)$ -dimensional space \mathcal{H} . Then we see that the intersection $\mathcal{H} \cap \mathcal{O}_i$ is a closed $(d - 1)$ -dimensional half-space of \mathcal{H} with \mathcal{I} as its boundary (noting again that \mathbf{o} lies on this boundary).

Thus each intersection $\mathcal{H} \cap \mathcal{O}_i$ is either a balanced $(d-2)$ -fan in \mathcal{H} whose axis passes through \mathbf{o} , a $(d-2)$ -dimensional hyperplane in \mathcal{H} passing through \mathbf{o} , a closed $(d-1)$ -dimensional half-space of \mathcal{H} whose boundary passes through \mathbf{o} or the entire space \mathcal{H} .

So let \mathcal{P} be the convex $(d-1)$ -dimensional polytope formed by intersecting our original polytope with \mathcal{H} . Since \mathbf{o} is interior to the original polytope and $\mathbf{o} \in \mathcal{H}$, it follows that \mathbf{o} is also interior to the new polytope \mathcal{P} . Furthermore, from our analysis of the intersections $\mathcal{H} \cap \mathcal{O}_i$ we see that we in fact have a new type I representation of our projective edge weight solution space expressed as the new polytope \mathcal{P} intersected with objects $\mathcal{H} \cap \mathcal{O}_1, \mathcal{H} \cap \mathcal{O}_2, \dots, \mathcal{H} \cap \mathcal{O}_m$ (where we ignore any objects for which $\mathcal{H} \cap \mathcal{O}_i = \mathcal{H}$ since these will not affect the final intersection).

Thus we have produced a new type I representation of our projective solution space whose enclosing polytope has smaller dimension $d-1$. Because dimension is bounded below by zero we can only repeat this procedure finitely many times, and so we must eventually reach a point at which our type I representation no longer contains any hyperplanes. Hence we have proven the following.

Type II Representation: The projective edge weight solution space is the intersection of $\leq d$ objects within a convex d -dimensional polytope, where $d \leq n$, \mathbf{o} is an interior point of this polytope and each of these objects is one of the following.

- A balanced $(d-1)$ -fan whose axis passes through \mathbf{o} ;
- A closed d -dimensional half-space whose boundary hyperplane passes through \mathbf{o} .

Finally, we observe that the intersection of one or more closed d -dimensional half-spaces with boundary passing through \mathbf{o} is in fact a convex cone whose vertex is \mathbf{o} (noting that a single half-space is itself such a cone). The result then follows. ■

What is particularly interesting about Theorem 5.5.1 is that the only step in this procedure that will increase the number of maximal faces of the solution space is where we take the intersection of $\leq d$ balanced $(d-1)$ -fans in d -dimensional space. Thus we can essentially ignore the variety of possible structures that the matching equations for an individual tetrahedron can form and simply focus our bounding arguments on this intersection of fans.

Lemma 5.5.2 Consider a d -dimensional polytope \mathcal{P} with interior point \mathbf{o} . Let \mathcal{S} be the space formed by the intersection of this polytope with some number of balanced $(d-1)$ -fans all of whose axes pass through \mathbf{o} . Then the following facts are true.

1. Each maximal face of \mathcal{S} can be expressed as the intersection of polytope \mathcal{P} and precisely one non-strict leaf from each fan.
2. Take any selection of strict leaves, one from each fan. If the intersection of these strict leaves in polytope \mathcal{P} is non-empty then the intersection of the corresponding non-strict leaves in polytope \mathcal{P} forms a maximal face of \mathcal{S} .

Proof Observe that a maximal face of \mathcal{S} is simply a maximal convex subset of \mathcal{S} , i.e., a convex subset of \mathcal{S} that is not a proper subset of some larger convex subset of \mathcal{S} . We prove each claim in turn.

1. Let \mathcal{F} be some maximal face of \mathcal{S} . Since \mathcal{F} is convex, it must be a convex subset of each fan. But any convex subset of a balanced fan must be contained entirely within one of its non-strict leaves, and so \mathcal{F} is a convex subset of some non-strict leaf from each fan.

Thus there is some intersection \mathcal{I} formed as the intersection of polytope \mathcal{P} and precisely one non-strict leaf from each fan for which $\mathcal{F} \subseteq \mathcal{I}$. But since each non-strict leaf is convex we see that \mathcal{I} is also a convex subset of \mathcal{S} , and so the maximality of \mathcal{F} implies that $\mathcal{F} = \mathcal{I}$.

2. We are given a selection of strict leaves $\mathcal{L}_1, \mathcal{L}_2, \dots$, one from each fan. Let \mathcal{F} represent the intersection of the corresponding non-strict leaves, and say that \mathcal{F} is not a maximal face of \mathcal{S} .

Since each non-strict leaf is convex it follows that \mathcal{F} is also convex, and so there must be some maximal face (i.e., maximally convex subset) $\mathcal{F}' \subseteq \mathcal{S}$ containing \mathcal{F} . From the previous result of this lemma we know that \mathcal{F}' can also be expressed as the intersection of polytope \mathcal{P} and precisely one non-strict leaf from each fan. Let $\mathcal{L}'_1, \mathcal{L}'_2, \dots$ denote the strict leaves corresponding to the non-strict leaves used to create \mathcal{F}' .

Since our original strict leaves $\mathcal{L}_1, \mathcal{L}_2, \dots$ have non-empty intersection, there is some point $\mathbf{p} \in \mathcal{F}$ belonging to each \mathcal{L}_i . But since $\mathcal{F} \subseteq \mathcal{F}'$, it follows that \mathbf{p} also belongs to each non-strict leaf used to create \mathcal{F}' . Furthermore, since \mathbf{p} lies on a strict leaf of each fan (and thus not the axis) we see that \mathbf{p} must in fact belong to strict leaves $\mathcal{L}'_1, \mathcal{L}'_2, \dots$. Finally, since the strict leaves of each fan are disjoint it follows that $\{\mathcal{L}_i\} = \{\mathcal{L}'_i\}$ and hence $\mathcal{F}' = \mathcal{F}$. So \mathcal{F} is indeed a maximal face of \mathcal{S} .

Thus each statement is proven. ■

Lemma 5.5.3 Consider a d -dimensional polytope \mathcal{P} with interior point \mathbf{o} and let \mathcal{S} be the space formed by the intersection of this polytope with k balanced $(d-1)$ -fans all of whose axes pass through \mathbf{o} .

Take any intersection of r axes and s non-strict leaves from these fans, where $r+s = k$ and these leaves and axes are each taken from distinct fans (so there is precisely one leaf or axis from each of the k available fans). Then the intersection of these r axes and s non-strict leaves in polytope \mathcal{P} forms a face of \mathcal{S} .

Proof Note that polytope \mathcal{P} is convex (and hence pseudo-convex) with the entirety of \mathcal{P} forming one of its faces. Recall also from Lemma 5.3.5 that each balanced fan is pseudo-convex with the axis and each non-strict leaf forming a face of the fan. It is then immediate from Lemma 5.1.6 that the intersection of \mathcal{P} with the selected r axes and s leaves (call this intersection \mathcal{I}) forms a face of the intersection of \mathcal{P} with the given k balanced fans, i.e., that \mathcal{I} forms a face of \mathcal{S} . ■

Definition 5.5.4 For each pair of integers d, k with $2 \leq d$ and $0 \leq k$, let $U(d, k)$ denote the smallest number with the following property. For any convex d -dimensional polytope, any point \mathbf{o} in the interior of this polytope and any set of precisely k balanced $(d-1)$ -fans in this polytope all of whose axes pass through \mathbf{o} , the complex formed by the intersection of these $(d-1)$ -fans within this polytope has at most $U(d, k)$ maximal faces.

Example 5.5.5 Figure 5.13 illustrates a scenario as described in Definition 5.5.4 with $d = 3$ and $k = 2$, although the enclosing polytope is not drawn. In this case we see five maximal faces in the intersection, each of which is an edge emanating from the central point \mathbf{o} .




Figure 5.13: A pair of balanced 2-fans intersecting within a 3-dimensional space

It can be seen by way of a detailed case analysis (which will not be reproduced here) that $U(3, 2) = 5$ and so the example illustrated has as many maximal faces as are possible for these values of d and k .

Lemma 5.5.6 For all pairs of integers d, k with $2 \leq d \leq d'$ and $0 \leq k \leq k'$, it is true that $U(d, k) \leq U(d', k')$.

Proof It is clear that $U(d, k) \leq U(d, k + 1)$ since any intersection of k fans can be expressed as an intersection of $k + 1$ fans by making two of the fans identical. The only case in which this is not possible is when $k = 0$, but in this case it is simple to see that $U(d, 0) = 1$ and $U(d, 1) = 3$ (this is proven explicitly in Lemmas 5.5.8 and 5.5.10).

We also see that $U(d, k) \leq U(d + 1, k)$ since any intersection of k balanced $(d - 1)$ -fans in a convex d -dimensional polytope can be converted into an intersection of k balanced d -fans in a convex $(d + 1)$ -dimensional polytope simply by taking the product of the original d -dimensional diagram with an interval I in the $(d + 1)$ th dimension (and declaring the new point \mathbf{o} to be some interior point of the product interval $\mathbf{o} \times I$).

Together these facts prove the required result. ■

Lemma 5.5.7 The projective edge weight solution space for any one-vertex n -tetrahedron triangulation where $n \geq 2$ has at most $U(n, n)$ maximal faces.

Proof Let \mathcal{S} be this projective edge weight solution space. From Theorem 5.5.1 we know that \mathcal{S} is formed from the intersection of a convex d -dimensional polytope with interior point \mathbf{o} , at most d balanced $(d - 1)$ -fans whose axes pass through \mathbf{o} and optionally a convex cone whose vertex is \mathbf{o} , for some $d \leq n$.

Consider first the intersection of just the polytope and the fans; call this partial intersection \mathcal{S}' . By definition of $U(d, k)$ we see that \mathcal{S}' has $\leq U(d, k)$ maximal faces, and by Lemma 5.5.6 we see that \mathcal{S}' has $\leq U(d, k) \leq U(d, d) \leq U(n, n)$ maximal faces. If $\mathcal{S} = \mathcal{S}'$ then we already have our desired result.

Otherwise the final solution space \mathcal{S} is created by intersecting \mathcal{S}' with some convex cone \mathcal{C} . Consider then any maximal face F of \mathcal{S} , and let F' be a maximal face of \mathcal{S}' containing F . Since \mathcal{C} is convex we see that $\mathcal{C} \cap F'$ is a face of \mathcal{S} containing F , and since F is maximal this face must be F itself.

Therefore each maximal face of \mathcal{S} is expressible as the intersection of \mathcal{C} with a maximal face of \mathcal{S}' , and so the number of maximal faces of \mathcal{S} is at most the number of maximal faces of \mathcal{S}' , which we already know to be at most $U(n, n)$. ■

Thus we devote our efforts to finding bounds for $U(d, k)$ since this gives us immediate bounds upon the number of maximal faces in the projective edge weight solution space.

We begin by dealing with the trivial case $U(d, 0)$ and then calculating a rough upper bound for all $U(d, k)$.

Lemma 5.5.8 For all $d \geq 2$, it is true that $U(d, 0) = 1$.

Proof The intersection of a convex polytope with no fans is merely the polytope itself. Thus the intersection has a single maximal face, this being the entire polytope. ■

Lemma 5.5.9 For all pairs of integers d, k satisfying the constraints of Definition 5.5.4, it is true that $U(d, k) \leq 3^k$.

Proof Let \mathcal{S} be the intersection of a polytope and k fans as described in Definition 5.5.4. From Lemma 5.5.2 we recall that each maximal face of \mathcal{S} can be expressed as the intersection of the entire polytope and precisely one leaf from each of these k fans. Since there are 3^k different ways of choosing one leaf from each fan, it follows that \mathcal{S} can have at most 3^k maximal faces. ■

In fact we can show that this upper bound is exact for small values of k .

Lemma 5.5.10 If d and k are integers for which $2 \leq d$ and $0 \leq 2k \leq d$ then $U(d, k) = 3^k$.

Proof Working in \mathbb{R}^d , let \mathbf{o} be the origin and let \mathcal{P} be some convex d -dimensional polytope with \mathbf{o} in its interior. For each $i \in \{1, 2, \dots, k\}$ define the $(d-1)$ -fan \mathcal{F}_i to be the union of the following half-hyperplanes.

- $x_{2i-1} = 0$ and $x_{2i} \geq 0$;
- $x_{2i-1} = x_{2i}$ and $x_{2i} \leq 0$;
- $x_{2i-1} = -x_{2i}$ and $x_{2i} \leq 0$.

We can observe that each \mathcal{F}_i is a balanced $(d-1)$ -fan with axis $x_{2i-1} = x_{2i} = 0$ and that each such axis contains \mathbf{o} . Furthermore, consider any selection of k strict leaves, one from each fan. Then it is simple to construct a point $\mathbf{v} \in \mathcal{P}$ belonging to all of these strict leaves by setting (v_{2i-1}, v_{2i}) to $(0, \varepsilon)$, $(-\varepsilon, -\varepsilon)$ or $(\varepsilon, -\varepsilon)$ according to which leaf of \mathcal{F}_i has been selected, where ε is some small positive constant.

Thus by Lemma 5.5.2 we see that for any selection of k non-strict leaves, one from each fan, the intersection of \mathcal{P} with all of these non-strict leaves forms a maximal face of the complete

intersection $\mathcal{P} \cap \mathcal{F}_1 \cap \dots \cap \mathcal{F}_k$. Since there are 3^k ways of making such a selection it follows that we have $\geq 3^k$ such maximal faces.

This construction therefore shows that $U(d, k) \geq 3^k$. But by Lemma 5.5.9 we know that $U(d, k) \leq 3^k$, and so for these values of d and k we must have $U(d, k) = 3^k$. \blacksquare

The bound $U(d, k) \leq 3^k$, though exact for small k as seen by Lemma 5.5.10, becomes quite rough as k approaches d . As an example, Table 5.1 lists the values of $U(d, k)$ for small d (as evaluated by hand using laborious case analyses) in comparison to this bound of 3^k .

(d, k)	(2, 0)	(2, 1)	(2, 2)	(3, 0)	(3, 1)	(3, 2)	(3, 3)
$U(d, k)$	1	3	3	1	3	5	5
3^k	1	3	9	1	3	9	27

(d, k)	(4, 0)	(4, 1)	(4, 2)	(4, 3)	(4, 4)
$U(d, k)$	1	3	9	9	9
3^k	1	3	9	27	81

Table 5.1: Comparisons of $U(d, k)$ with the rough bound 3^k for small d

As well as the upper bound described in Lemma 5.5.9, we can also give a general lower bound for $U(d, k)$.

Lemma 5.5.11 For all pairs of integers d, k with $2 \leq d$ and $1 \leq k \leq d - 1$, it is true that $2^k + 1 \leq U(d, k)$.

Proof Consider the case in which $k = d - 1$. Working in \mathbb{R}^d , let \mathbf{o} be the origin and let \mathcal{P} be some convex d -dimensional polytope for which \mathbf{o} is an interior point. For each $i \in \{1, 2, \dots, d - 1\}$ define the $(d - 1)$ -fan \mathcal{F}_i to be the union of the following half-hyperplanes.

- $x_i = 0$ and $x_d \geq 0$;
- $x_i = x_d$ and $x_d \leq 0$;
- $x_i = -x_d$ and $x_d \leq 0$.

It can be seen that each \mathcal{F}_i is a balanced $(d - 1)$ -fan with axis $x_i = x_d = 0$ which passes through \mathbf{o} . The intersection of these fans is the set of all vectors of the form

$$(\pm\lambda, \pm\lambda, \dots, \pm\lambda, -\lambda) \quad \text{or} \quad (0, 0, \dots, 0, \lambda) \quad (5.4)$$

for $\lambda \geq 0$. This is a union of disjoint and non-parallel rays leaving \mathbf{o} (illustrated for $d = 3$ in Figure 5.14). There are $2^{d-1} + 1$ such rays corresponding to the 2^{d-1} choices of $+$ or $-$ for vectors of the first form and the single ray of vectors of the second form.

Thus, when this set of fans is intersected with polytope \mathcal{P} , the set of maximal faces will consist of $2^{d-1} + 1 = 2^k + 1$ edges, one for each of these rays. Since our scenario satisfies the requirements of Definition 5.5.4 it follows that $2^k + 1 \leq U(d, k)$.

Hence our result is established when $k = d - 1$. If $k < d - 1$ then our earlier argument gives $2^k + 1 \leq U(k + 1, k)$, from which Lemma 5.5.6 gives $2^k + 1 \leq U(k + 1, k) \leq U(d, k)$. \blacksquare




Figure 5.14: The intersection described by Equation 5.4 in the case where $d = 3$

Example 5.5.12 Example 5.5.5, illustrated earlier in Figure 5.13, shows precisely how the construction used to prove Lemma 5.5.11 appears in the case where $d = 3$ and $k = 2$.

Corollary 5.5.13 For all pairs of integers d, k with $2 \leq d$ and $k \geq d - 1$, it is true that $2^{d-1} + 1 \leq U(d, k)$.

Proof From Lemma 5.5.11 we have $2^{d-1} + 1 \leq U(d, d - 1)$, and from Lemma 5.5.6 we have $U(d, d - 1) \leq U(d, k)$. ■

Now that we have both an upper and a lower bound for $U(d, k)$, we are faced with the question of where between these bounds the actual values for $U(d, k)$ lie. We have already established that the upper bound of 3^k is exact for small k . It is the author's belief that the lower bound expressed in Corollary 5.5.13 for large k is also exact, as conjectured below.

Conjecture 5.5.14 For all pairs of integers d, k with $2 \leq d$ and $k \geq d - 1$, we have $U(d, k) = 2^{d-1} + 1$. In particular, $U(d, d - 1) = U(d, d) = 2^{d-1} + 1$.

The conjecture that $U(d, d - 1) = 2^{d-1} + 1$ was initially formed by observing this fact for very small d and then constructing scenarios for all d in which the intersection of $d - 1$ balanced $(d - 1)$ -fans has precisely $2^{d-1} + 1$ maximal faces, as described in Lemma 5.5.11. This belief was further encouraged when the expression $2^{d-1} + 1$ came naturally out of the algebra in the case where the fans intersect in general position, as described in Chapter 6.

Intuitively it feels that placing the fans in general position offers the opportunity for the greatest number of maximal faces, since a slight perturbation of a leaf-leaf intersection will merely result in a smaller leaf-leaf intersection, whereas a slight perturbation of a leaf-axis intersection or an axis-axis intersection can result in a union of several leaf-leaf intersections (since each axis forms the boundary of three leaves), potentially increasing the maximal face count as a result. For very small d it has indeed been observed that the maximal number of faces is only realised when the fans are placed in general position.

The conjecture that $U(d, k) = U(d, d - 1)$ for $k \geq d - 1$ was again based upon experimental results for very small d . However, if we observe how the maximal face count grows as we intersect one fan after another, we see that as we add a new fan we can break a maximal face into two or three maximal faces by having it meet two or three of the leaves of this new fan. Each such splitting produces maximal faces of smaller dimension, and so this operation cannot continue indefinitely.

When working in a d -dimensional space, we see that the intersection of $d - 1$ hyperplanes of dimension $d - 1$ in general position is a line, the smallest possible dimension for a non-trivial face

(since all faces include the central point \mathbf{o}). Thus, combining this with the previous observation, it makes intuitive sense that we should not expect more maximal faces by intersecting more than $d - 1$ fans in a d -dimensional space. We therefore expect that the bound $U(d, k)$ should cease to grow once $k \geq d - 1$.

If Conjecture 5.5.14 is indeed true, Lemma 5.5.7 then implies the following result, which is the primary motivation behind this and the following chapter.

Consequence 5.5.15 (Maximal Face Upper Bound) Let T be a triangulation of a closed 3-manifold with one vertex and n tetrahedra. Then the projective edge weight solution space for T has at most $2^{n-1} + 1$ maximal faces. Equivalently (by Theorem 5.4.7), the projective standard solution space for T has at most $2^{n-1} + 1$ maximal embedded faces.

With Consequence 5.5.15 as our ultimate goal, we aim to investigate Conjecture 5.5.14 using the geometric tools developed within this chapter. We proceed then to Chapter 6, in which these geometric tools are used in conjunction with a combinatorial analysis for the general position scenario outlined above.

Chapter 6

General Position Bounds

In this chapter we continue the investigation begun in Chapter 5 into the complexity of the projective solution space. Consider once more some intersection of k balanced $(d-1)$ -fans in a convex d -dimensional polytope. As discussed in the closing remarks of Chapter 5, it seems intuitive that such an intersection will have the largest possible number of maximal faces (i.e., $U(d, k)$ maximal faces) in a case where these fans are placed in general position.

Thus we continue our pursuit of tighter bounds for $U(d, k)$ by examining this general position scenario. Restricting ourselves to fans in general position also makes our task somewhat easier by removing a myriad of special cases that would otherwise need to be considered. Once we understand this general position scenario we can then form a better idea of how we might tackle the larger unrestricted problem.

6.1 Definitions

We begin with a formal discussion of general position and the presentation of some core results.

Definition 6.1.1 (General Position) A collection of balanced $(d-1)$ -fans in a convex d -dimensional space \mathcal{P} , all of whose axes pass through some point $\mathbf{o} \in \text{int}(\mathcal{P})$, is in *general position* if the following conditions hold.

- For any $r \geq 1$, the intersection of any r axes of different fans is a $(d-2r)$ -dimensional hyperplane in \mathcal{P} containing \mathbf{o} in its relative interior. If $d-2r \leq 0$ then this intersection is simply \mathbf{o} itself.
- For any $s \geq 1$, if the intersection of some s strict leaves from different fans is non-empty then this intersection is in fact an open $(d-s)$ -dimensional cone (i.e., excluding its boundary). If $d-s \leq 0$ then such an intersection must be empty.
- Furthermore, assume $s \geq 1$ and the intersection of some s strict leaves is non-empty as described above. Call the open cone formed by this intersection \mathcal{C} , and call its closure (formed by the intersection of the corresponding s non-strict leaves) $\bar{\mathcal{C}}$.

Consider now any additional r axes for which the above s leaves and these r axes all come from $r + s$ distinct fans. Let the intersection of these r axes be \mathcal{A} (noting from above that \mathcal{A} is a $(d - 2r)$ -dimensional hyperplane). Then one of the following cases must occur.

- Hyperplane \mathcal{A} does not meet the open cone \mathcal{C} at all (this may occur for instance when \mathcal{A} lies entirely to one side of the cone \mathcal{C}). In this case, the overall intersection $\bar{\mathcal{C}} \cap \mathcal{A}$ of non-strict leaves and axes must be the single point \mathbf{o} . That is, the hyperplane \mathcal{A} may not meet any part of the boundary of cone $\bar{\mathcal{C}}$ except for its vertex \mathbf{o} .
- Hyperplane \mathcal{A} meets the open cone \mathcal{C} . In this case \mathcal{A} forms a hyperplane through the open cone \mathcal{C} so that the overall intersection $\bar{\mathcal{C}} \cap \mathcal{A}$ of non-strict leaves and axes forms a closed $(d - 2r - s)$ -dimensional cone. If $d - 2r - s \leq 0$ then this case cannot occur.

Although Definition 6.1.1 seems restrictive, it merely begins with the general position concept of k different $(d - 1)$ -dimensional hyperplanes intersecting in a $(d - k)$ -dimensional hyperplane and adapts this concept for our particular scenario, in which some hyperplanes are only half-hyperplanes (the leaves of the fans) and in which all of our objects pass through the common point \mathbf{o} . In essence, this definition attempts to remain true to the idea that “general position” should describe precisely those scenarios in which a small perturbation of the fan positions does not change the combinatorial structure of their intersections.

Example 6.1.2 The fans described in Example 5.5.5 (illustrated in Figure 5.13 on page 182) are in general position. Likewise, the sets of fans constructed in the proof of Lemma 5.5.10 are in general position.




Figure 6.1: Two fans in \mathbb{R}^3 that are not in general position

The fans illustrated in Figure 6.1 (for which $d = 3$) are not in general position since an axis and a leaf can be found that intersect in a 1-dimensional ray even though $d - 2r - s = 0$ when $r = s = 1$.

Lemma 6.1.3 Consider a collection of k balanced $(d - 1)$ -fans in general position in a convex d -dimensional space \mathcal{P} , all of whose axes pass through some point $\mathbf{o} \in \text{int}(\mathcal{P})$. Furthermore, consider some k non-strict leaves $\mathcal{L}_1, \dots, \mathcal{L}_k$, one from each fan.

Let \mathcal{I} be the intersection $\mathcal{L}_1 \cap \dots \cap \mathcal{L}_k$. If \mathcal{I} contains at least one point from each of the corresponding k strict leaves, then \mathcal{I} contains at least one point \mathbf{p} for which \mathbf{p} simultaneously belongs to all k corresponding strict leaves.

Proof For each non-strict leaf \mathcal{L}_i , let \mathcal{L}'_i denote the corresponding strict leaf. Furthermore, say \mathcal{I} contains points $\mathbf{p}_1, \dots, \mathbf{p}_k$ where $\mathbf{p}_i \in \mathcal{L}'_i$ for each i . We construct the new point

$$\mathbf{p} = \frac{1}{n}(\mathbf{p}_1 + \dots + \mathbf{p}_k),$$

and claim that \mathbf{p} simultaneously belongs to each of the strict leaves \mathcal{L}'_i .

First note that since \mathcal{I} is the intersection of non-strict leaves \mathcal{L}_i , it is true that $\mathbf{p}_i \in \mathcal{L}_j$ for each i, j . Let \mathcal{C} denote the convex hull of points $\mathbf{p}_1, \dots, \mathbf{p}_k$; then we also have $\mathcal{C} \subseteq \mathcal{L}_i$ for each i . Since \mathbf{p} is a convex combination of the \mathbf{p}_i it also follows that $\mathbf{p} \in \mathcal{L}_i$ for each i .

Now consider any non-strict leaf \mathcal{L}_i . Define point

$$\mathbf{p}'_i = \frac{1}{n-1}(\mathbf{p}_1 + \dots + \mathbf{p}_{i-1} + \mathbf{p}_{i+1} + \dots + \mathbf{p}_k),$$

where the sum inside the brackets includes every point \mathbf{p}_j for which $i \neq j$. The new point \mathbf{p}'_i is also in the convex hull \mathcal{C} and so $\mathbf{p}'_i \in \mathcal{L}_i$. Furthermore, we see that $\mathbf{p} = \frac{1}{n}\mathbf{p}_i + \frac{n-1}{n}\mathbf{p}'_i$. Consider the following two cases.

- It might be true that $\mathbf{p}_i = \mathbf{p}'_i$. In this case we have $\mathbf{p} = \mathbf{p}_i$ and so $\mathbf{p} \in \mathcal{L}'_i$.
- Otherwise $\mathbf{p}_i \neq \mathbf{p}'_i$, in which case we see that \mathbf{p} lies in the interior of the line segment joining \mathbf{p}_i and \mathbf{p}'_i . In this case, since $\mathbf{p}_i, \mathbf{p}'_i \in \mathcal{L}_i$ and $\mathbf{p}_i \in \text{relint}(\mathcal{L}_i) = \mathcal{L}'_i$, it follows by the convexity of \mathcal{L}_i that we also have $\mathbf{p} \in \text{relint}(\mathcal{L}_i) = \mathcal{L}'_i$.

So either way we have $\mathbf{p} \in \mathcal{L}'_i$. Thus we see that \mathbf{p} belongs to all k strict leaves as required. \blacksquare

Lemma 6.1.4 Consider a collection of k balanced $(d-1)$ -fans in general position within a convex d -dimensional space \mathcal{P} , all of whose axes pass through some point $\mathbf{o} \in \text{int}(\mathcal{P})$. If $k \geq d$ then the intersection of these fans is the single point \mathbf{o} .

Proof Take any point \mathbf{p} belonging to this intersection. Since \mathbf{p} belongs to each fan, let \mathbf{p} belong to the axes of r of these fans and to strict leaves of s of these fans, where $r + s = k$.

- If $s = 0$ then $r = k$ and $d - 2r = d - 2k \leq 0$. So Definition 6.1.1 implies that the intersection of these r axes is the single point \mathbf{o} , and so $\mathbf{p} = \mathbf{o}$.
- If $r = 0$ then $s = k$ and $d - s = d - k \leq 0$. So Definition 6.1.1 implies that \mathbf{p} cannot belong to all s of these strict leaves.
- Otherwise $r, s > 0$. But then $d - 2r - s < d - (r + s) = d - k \leq 0$, and so Definition 6.1.1 implies that the intersection of these r axes cannot meet the intersection of these s strict leaves.

Thus $s = 0$ and we see that in fact $\mathbf{p} = \mathbf{o}$, so our overall intersection of fans can be at most \mathbf{o} . Finally observe that this overall intersection must contain \mathbf{o} since each fan passes through \mathbf{o} . \blacksquare

Lemma 6.1.4 is in line with the remarks following Conjecture 5.5.14, in which we expect $U(d, k)$ to stop growing for $k \geq d$. Thus, for these general position arguments, we focus our attention on the intersections of at most $d-1$ fans within a d -dimensional space.

6.2 Intersection Pieces

Recall from Lemma 5.5.7 that for the purpose of counting maximal faces our projective edge weight solution space can be viewed as an intersection of k balanced fans within a d -dimensional polytope. In order to bound this number of maximal faces in this general position scenario, our focus lies upon bounding the number of intersections of k strict leaves, as described in Theorem 6.2.8.

However, it is difficult to find an immediate tight bound upon this quantity. Instead we find it easier to bound the number of regions into which our polytope and various structures within it are divided by the intersecting fans, such as in Lemma 6.3.1 where we prove that the region count is $\leq 3^k$.

Thus our strategy is to examine the way in which these fans decompose the polytope into cells of various dimensions. Using a multitude of Euler characteristic equations we can then convert our bounds upon various region counts into the desired bound upon the number of leaf intersections.

Definitions 6.2.1 (Pieces and Piece Counts) Consider a collection of k balanced $(d-1)$ -fans in general position in a convex d -dimensional space \mathcal{P} , all of whose axes pass through some point $\mathbf{o} \in \text{int}(\mathcal{P})$, where $1 \leq k \leq d-1$. We examine here a variety of combinatorial properties and components of this structure.

- The number of maximal faces of the intersection of all k fans within \mathcal{P} is called the *maximal face count* and denoted by Φ .
- Since each fan divides \mathcal{P} into three d -dimensional regions, the union of all k fans likewise divides \mathcal{P} into some number of d -dimensional regions. Call this number of regions the *region count* and denote it by R .
- Consider any $r, s \geq 0$. Suppose that some collection of s strict leaves and r axes all from $r+s$ distinct fans has non-empty intersection, and that this intersection is at least 1-dimensional (i.e., is not just the single point \mathbf{o}). Then we call the intersection of the corresponding s non-strict and r axes an *original (r, s) -piece*.

If $s = 0$ then the original (r, s) -pieces are simply the intersections of r axes from distinct fans for which this intersection is at least 1-dimensional. The only original $(0, 0)$ -piece is the entire polytope \mathcal{P} .

- Consider now any original (r, s) -piece \mathcal{O} . Recall from above that the entire set of fans divides \mathcal{P} into a smaller number of d -dimensional regions. In a similar fashion, the remaining $k-r-s$ fans that are not used in creating the original piece \mathcal{O} divide \mathcal{O} into some smaller number of regions each of the same dimension as \mathcal{O} . Each of these smaller regions is called a *final (r, s) -piece*.

Final (r, s) -pieces include their boundaries, i.e., they are closed. Note that there are R final $(0, 0)$ -pieces representing the regions into which \mathcal{P} is subdivided by the entire set of fans.

- For each r and s , denote the total number of original (r, s) -pieces by $G_{r,s}$ and the total number of final (r, s) -pieces by $F_{r,s}$. $G_{r,s}$ is called the *original (r, s) -piece count* and $F_{r,s}$ is called the *final (r, s) -piece count*. If $r+s > k$ then we define $G_{r,s} = F_{r,s} = 0$.

Example 6.2.2 Consider the pair of 2-fans in 3-dimensional space, illustrated in Figure 6.2, where the surrounding polytope has been omitted for the clarity of the diagram.




Figure 6.2: Original and final pieces for a pair of 2-fans

As noted in Definitions 6.2.1 there is a single original $(0, 0)$ -piece, this being the polytope itself. The polytope is subdivided by the two fans into nine 3-dimensional regions, leaving $R = 9$ final $(0, 0)$ -pieces.

There are two axes (one from each fan) giving two original $(1, 0)$ -pieces. Each axis is subdivided into two pieces by the other fan at the central point \mathbf{o} , leaving four final $(1, 0)$ -pieces. There are six original $(0, 1)$ -pieces (three leaves from each fan), which are subdivided into 16 different two-dimensional discs, leaving 16 final $(0, 1)$ -pieces.

Examining the intersections between strict leaves we see five leaf-leaf intersections in total (one ray pointing up and four rays pointing down). Thus we have five original $(0, 2)$ -pieces, none of which are subdivided any further, leaving also five final $(0, 2)$ -pieces. In addition, the intersection of our two fans consists entirely of these five rays with their boundary points and so each ray forms a maximal face of this intersection giving maximal face count $\Phi = 5$.

Each leaf-axis or axis-axis intersection is simply the single point \mathbf{o} . Since these are all 0-dimensional structures, we have no $(1, 1)$ -pieces or $(2, 0)$ -pieces.

These results are summarised in Table 6.1.

$G_{r,s}$	$s = 0$	$s = 1$	$s = 2$	$F_{r,s}$	$s = 0$	$s = 1$	$s = 2$
$r = 0$	1	6	5	$r = 0$	9	16	5
$r = 1$	2	0	0	$r = 1$	4	0	0
$r = 2$	0	0	0	$r = 2$	0	0	0

Table 6.1: Original and final piece counts for Figure 6.2

Lemma 6.2.3 Consider the structures formed by the original and final (r, s) -pieces described in Definitions 6.2.1.

- If $2r + s \geq d$ then there are no original or final (r, s) -pieces.
- If $2r + s < d$ then each original and final (r, s) -piece can be considered as a convex $(d - 2r - s)$ -dimensional cone projecting from vertex \mathbf{o} (though note that the entire polytope \mathcal{P} is itself such a cone).

Proof Suppose $2r + s \geq d$. If $s = 0$ then an original (r, s) -piece is an intersection of r axes where $2r \geq d$, but general position implies that this intersection can only be the single point \mathbf{o} . If $s > 0$ then because $2r + s \geq d$ general position implies that the intersection of s strict leaves cannot meet the intersection of r axes from different fans. Either way we see that we can have no original (r, s) -pieces, and since final pieces are simply subdivisions of original pieces we can have no final (r, s) -pieces either.

So let $2r + s < d$. Since every non-strict leaf and every axis can be seen as a convex cone with vertex \mathbf{o} (noting that any hyperplane through \mathbf{o} is itself such a cone), any intersection of non-strict leaves and axes (and thus any original piece) can also be seen as a convex cone with vertex \mathbf{o} .

Furthermore, if $s = 0$ then an original (r, s) -piece is an intersection of r axes which by general position is a $(d - 2r)$ -dimensional hyperplane in \mathcal{P} . If $s > 0$ then the s strict leaves involved in forming the piece must meet the r axes involved and so general position implies that the s non-strict leaves and r axes meet in a closed $(d - 2r - s)$ -dimensional cone. Either way we see that each original (r, s) -piece is a convex $(d - 2r - s)$ -dimensional cone with vertex \mathbf{o} .

Each final (r, s) -piece \mathcal{F} is created by the subdivision of an original (r, s) -piece \mathcal{O} by the remaining fans. Since each fan divides polytope \mathcal{P} into three regions whose closures are convex cones with vertex \mathbf{o} , we can view \mathcal{F} as the intersection of \mathcal{O} with a collection of such closed regions and so \mathcal{F} is also a convex cone with vertex \mathbf{o} . Furthermore, the dimension of \mathcal{F} is the same as the dimension of \mathcal{O} , i.e., \mathcal{F} is $(d - 2r - s)$ -dimensional. \blacksquare

Lemma 6.2.4 In the scenario described in Definitions 6.2.1, any intersection of r axes and s non-strict leaves from $r + s$ distinct fans is either the single point \mathbf{o} or an original (r', s') -piece where $r' + s' = r + s$ and $r' \geq r$.

Proof Let \mathcal{O} denote this intersection. If any non-strict leaf fails to meet \mathcal{O} in its leaf interior, we can replace this non-strict leaf with its boundary, i.e., the corresponding axis, without changing the intersection \mathcal{O} . Thus \mathcal{O} can be expressed as the intersection

$$\mathcal{O} = \bar{\mathcal{L}}_1 \cap \dots \cap \bar{\mathcal{L}}_{s'} \cap \mathcal{A}_{s'+1} \cap \dots \cap \mathcal{A}_{r'+s'},$$

where each $\bar{\mathcal{L}}_i$ is a non-strict leaf, each \mathcal{A}_i is an axis, each of the r' axes and s' leaves come from $r' + s'$ distinct fans, $r' + s' = r + s$ and $r' \geq r$. Furthermore, this can be done so that for each $\bar{\mathcal{L}}_i$ ($1 \leq i \leq s'$) our intersection \mathcal{O} contains at least one point from the relative interior of leaf $\bar{\mathcal{L}}_i$, i.e., from the corresponding strict leaf.

If $s' = 0$ then \mathcal{O} is simply the intersection of r' axes from different fans. If this intersection is at least 1-dimensional then from Definitions 6.2.1 it is an original (r', s') -piece; if not then it is the point \mathbf{o} .

If $s' > 0$ then we can apply Lemma 6.1.3 to the intersection of the s' fans from which our leaves are taken to show that the intersection $\bar{\mathcal{L}}_1 \cap \dots \cap \bar{\mathcal{L}}_{s'}$ contains some point simultaneously belonging to all s' corresponding strict leaves. Thus these s' strict leaves have non-empty intersection, and general position implies that these strict leaves intersect in some open $(d - s')$ -dimensional cone \mathcal{C} .

Again by general position we see that if the intersection $\mathcal{A}_{s'+1} \cap \dots \cap \mathcal{A}_{r'+s'}$ does not meet open cone \mathcal{C} then the intersection \mathcal{O} of non-strict leaves and axes must be the single point \mathbf{o} . Otherwise

this intersection of axes meets the intersection of strict leaves \mathcal{C} and from Definitions 6.2.1 we see that \mathcal{O} is an original (r', s') -piece. ■

Lemma 6.2.5 Consider the situation described in Definitions 6.2.1. Let original piece \mathcal{O} be formed from axes $\mathcal{A}_1, \dots, \mathcal{A}_r$ and strict leaves $\mathcal{L}_1, \dots, \mathcal{L}_s$ and let original piece \mathcal{O}' be formed from axes $\mathcal{A}'_1, \dots, \mathcal{A}'_{r'}$ and strict leaves $\mathcal{L}'_1, \dots, \mathcal{L}'_{s'}$.

If each axis \mathcal{A}_i appears in the list $\mathcal{A}'_1, \dots, \mathcal{A}'_{r'}$, and if each leaf \mathcal{L}_i either appears in the list $\mathcal{L}'_1, \dots, \mathcal{L}'_{s'}$ or has its corresponding axis in the list $\mathcal{A}'_1, \dots, \mathcal{A}'_{r'}$, then $\mathcal{O}' \subseteq \mathcal{O}$. Otherwise the intersection $\mathcal{O}' \cap \mathcal{O}$ has strictly smaller dimension than \mathcal{O}' and in particular $\mathcal{O}' \not\subseteq \mathcal{O}$.

Proof For each strict leaf \mathcal{L}_i or \mathcal{L}'_i let $\bar{\mathcal{L}}_i$ and $\bar{\mathcal{L}}'_i$ denote the closure, i.e., the corresponding non-strict leaf. Consider the intersection

$$\mathcal{O}' \cap \mathcal{O} = \mathcal{A}'_1 \cap \dots \cap \mathcal{A}'_{r'} \cap \mathcal{A}_1 \cap \dots \cap \mathcal{A}_r \cap \bar{\mathcal{L}}'_1 \cap \dots \cap \bar{\mathcal{L}}'_{s'} \cap \bar{\mathcal{L}}_1 \cap \dots \cap \bar{\mathcal{L}}_s. \quad (6.1)$$

If the conditions described in this lemma are true then each \mathcal{A}_i already appears as some \mathcal{A}'_j elsewhere in this equation and so can be safely removed without changing the overall intersection. Similarly if a leaf \mathcal{L}_i appears in the list $\mathcal{L}'_1, \dots, \mathcal{L}'_{s'}$, then it can also be removed from this equation. If a leaf \mathcal{L}_i has its corresponding axis in the list $\mathcal{A}'_1, \dots, \mathcal{A}'_{r'}$, then it too can be removed, since a non-strict leaf intersected with its corresponding axis is just this axis. Thus Equation 6.1 reduces to

$$\mathcal{O}' \cap \mathcal{O} = \mathcal{A}'_1 \cap \dots \cap \mathcal{A}'_{r'} \cap \bar{\mathcal{L}}'_1 \cap \dots \cap \bar{\mathcal{L}}'_{s'} = \mathcal{O}',$$

and so $\mathcal{O}' \subseteq \mathcal{O}$.

Otherwise the conditions described are not satisfied and so one of the following scenarios must occur. In each case we will examine ways of restricting the terms of Equation 6.1 to just a selection of the original terms, creating a new intersection that is a superset of $\mathcal{O}' \cap \mathcal{O}$.

- Some axis \mathcal{A}_i is not in the list $\mathcal{A}'_1, \dots, \mathcal{A}'_{r'}$. In this case we can restrict the terms of Equation 6.1 to form an intersection of $r' + 1$ axes and $s' - 1$ leaves all from distinct fans, which from Lemma 6.2.4 is an original $(r' + 1 + q, s' - 1 - q)$ -piece for some $q \geq 0$. From Lemma 6.2.3 this piece is $(d - 2r' - s' - 1 - q)$ -dimensional, observing that $d - 2r' - s' - 1 - q < d - 2r' - s' = t'$.
Note that if $s' = 0$ then we can just restrict our terms to an intersection of $r' + 1$ axes which is $(d - 2r' - 2)$ -dimensional where $d - 2r' - 2 < d - 2r' = t'$.
- $\{\mathcal{A}_1, \dots, \mathcal{A}_r\} \subseteq \{\mathcal{A}'_1, \dots, \mathcal{A}'_{r'}\}$ but some leaf \mathcal{L}_i does not come from any of the fans used in creating \mathcal{O}' . Here we can restrict the terms of Equation 6.1 to an intersection of r' axes and $s' + 1$ leaves all from distinct fans giving an original $(r' + q, s' + 1 - q)$ -piece for some $q \geq 0$. This piece is again $(d - 2r' - s' - 1 - q)$ -dimensional where $d - 2r' - s' - 1 - q < t'$.
- $\{\mathcal{A}_1, \dots, \mathcal{A}_r\} \subseteq \{\mathcal{A}'_1, \dots, \mathcal{A}'_{r'}\}$ but two leaves \mathcal{L}_i and \mathcal{L}'_j are distinct leaves from the same fan. In this case we can replace the intersection $\mathcal{L}_i \cap \mathcal{L}'_j$ with their common axis and thus restrict the terms of Equation 6.1 to an intersection of $r' + 1$ axes and $s' - 1$ leaves all from distinct fans which is again a $(d - 2r' - s' - 1 - q)$ -dimensional structure where $d - 2r' - s' - 1 - q < t'$.

We see then that in each case we can construct a superset of $\mathcal{O}' \cap \mathcal{O}$ of dimension $< t$. Thus the dimension of $\mathcal{O}' \cap \mathcal{O}$ is strictly less than t , and in particular $\mathcal{O}' \cap \mathcal{O} \neq \mathcal{O}'$ and so $\mathcal{O}' \not\subseteq \mathcal{O}$. ■

Corollary 6.2.6 In the situation described in Definitions 6.2.1, let original piece \mathcal{O} have dimension t and let original piece \mathcal{O}' have dimension t' . Then if $\mathcal{O} \neq \mathcal{O}'$, the intersection $\mathcal{O} \cap \mathcal{O}'$ has dimension strictly less than $\max(t, t')$.

Proof If $\mathcal{O} \neq \mathcal{O}'$ then either $\mathcal{O} \not\subseteq \mathcal{O}'$ or $\mathcal{O}' \not\subseteq \mathcal{O}$. From Lemma 6.2.5 it follows that the dimension of $\mathcal{O} \cap \mathcal{O}'$ is either $< t$ or $< t'$ and hence is $< \max(t, t')$. ■

Corollary 6.2.7 In the situation described in Definitions 6.2.1, any two original pieces formed from different collections of axes and leaves (as described in Definitions 6.2.1) must be distinct. Similarly, any two final pieces formed from different collections of axes and leaves must be distinct.

Proof Consider original piece \mathcal{O} formed from axes $\mathcal{A}_1, \dots, \mathcal{A}_r$ and strict leaves $\mathcal{L}_1, \dots, \mathcal{L}_s$ and original piece \mathcal{O}' formed from axes $\mathcal{A}'_1, \dots, \mathcal{A}'_{r'}$ and strict leaves $\mathcal{L}'_1, \dots, \mathcal{L}'_{s'}$.

If $\mathcal{O} = \mathcal{O}'$ then we have both $\mathcal{O} \subseteq \mathcal{O}'$ and $\mathcal{O}' \subseteq \mathcal{O}$. Applying Lemma 6.2.5 in both directions we see that $\{\mathcal{A}_1, \dots, \mathcal{A}_r\} = \{\mathcal{A}'_1, \dots, \mathcal{A}'_{r'}\}$, and armed with this information we can apply Lemma 6.2.5 again to show that $\{\mathcal{L}_1, \dots, \mathcal{L}_s\} = \{\mathcal{L}'_1, \dots, \mathcal{L}'_{s'}\}$. Hence \mathcal{O} and \mathcal{O}' are formed from the same collections of axes and leaves.

Consider now final pieces \mathcal{F} and \mathcal{F}' formed from the subdivisions of some original pieces \mathcal{O} and \mathcal{O}' respectively. Suppose that $\mathcal{F} = \mathcal{F}'$ but that these two pieces are formed from different collections of axes and leaves. Then both final pieces have the same dimension; let this be t . Both \mathcal{O} and \mathcal{O}' must therefore also have dimension t .

Since \mathcal{O} and \mathcal{O}' are formed from different collections of axes and leaves the above argument shows that $\mathcal{O} \neq \mathcal{O}'$, and from Corollary 6.2.6 the intersection $\mathcal{O} \cap \mathcal{O}'$ thus has dimension $< t$. But $\mathcal{O} \cap \mathcal{O}'$ includes the common final piece $\mathcal{F} = \mathcal{F}'$ which is t -dimensional, leading us to a contradiction. ■

Theorem 6.2.8 Consider a collection of k balanced $(d-1)$ -fans in general position in a convex d -dimensional space \mathcal{P} , all of whose axes pass through some point $\mathbf{o} \in \text{int}(\mathcal{P})$, where $1 \leq k \leq d-1$. If the intersection of these k fans within \mathcal{P} has any non-trivial maximal faces (i.e., contains more than just the single point \mathbf{o}), then the maximal faces are precisely the original $(r, k-r)$ -pieces \mathcal{O} for which either $r=0$ or $\mathcal{O} \not\subseteq \mathcal{O}'$ for any original $(r-1, k-(r-1))$ -piece \mathcal{O}' .

In particular if k is as large as possible, i.e., $k = d-1$, then these maximal faces are precisely the original $(0, k)$ -pieces.

Proof Assume that we have at least one non-trivial maximal face, and consider any non-trivial maximal face \mathcal{F} . Using Lemma 5.5.2 we see that \mathcal{F} can be expressed as the intersection of k non-strict leaves $\mathcal{L}_1, \dots, \mathcal{L}_k$, one from each fan. From Lemma 6.2.4 it follows that \mathcal{F} is an original (r, s) -piece where $r+s=k$.

Observe moreover that \mathcal{F} contains point \mathbf{o} and so \mathbf{o} itself cannot be a maximal face. Thus every maximal face is non-trivial and is therefore an original $(r, k-r)$ -piece for some r .

Conversely, let \mathcal{O} be some original $(r, k - r)$ -piece. If $\mathcal{O} \subseteq \mathcal{O}'$ for some different original $(r', k - r')$ -piece \mathcal{O}' then Corollary 6.2.7 shows that \mathcal{O} is a proper subset of \mathcal{O}' and Lemma 5.5.3 implies that \mathcal{O}' is also a face of the overall intersection of fans; thus \mathcal{O} cannot be a maximal face.

On the other hand, if \mathcal{O} is not a maximal face then there is some maximal face \mathcal{O}' for which $\mathcal{O} \subseteq \mathcal{O}'$ and by our earlier argument \mathcal{O}' must be an original $(r', k - r')$ -piece for some r' . Thus the maximal faces are precisely the original $(r, k - r)$ -pieces \mathcal{O} for which $\mathcal{O} \not\subseteq \mathcal{O}'$ for any different original $(r', k - r')$ -piece \mathcal{O}' .

Consider now some original $(r, k - r)$ -piece \mathcal{O} and some different original $(r', k - r')$ -piece \mathcal{O}' for which $\mathcal{O} \subseteq \mathcal{O}'$. Since axes and leaves from all k fans are used in creating both \mathcal{O} and \mathcal{O}' , the only way in which the conditions of Lemma 6.2.5 can be satisfied is to have $r' < r$ along with the following additional constraints.

- Piece \mathcal{O} is constructed from axes $\mathcal{A}_1, \dots, \mathcal{A}_{r'}, \mathcal{A}_{r'+1}, \dots, \mathcal{A}_r$ and strict leaves $\mathcal{L}_{r+1}, \dots, \mathcal{L}_k$.
- Piece \mathcal{O}' is constructed from axes $\mathcal{A}_1, \dots, \mathcal{A}_{r'}$ and strict leaves $\mathcal{L}_{r'+1}, \dots, \mathcal{L}_r, \mathcal{L}_{r+1}, \dots, \mathcal{L}_k$, where axis \mathcal{A}_i is the axis corresponding to leaf \mathcal{L}_i for each i in the range $r' < i \leq r$.

Furthermore we can consider the intersection of axes $\mathcal{A}_1, \dots, \mathcal{A}_{r'}, \mathcal{A}_{r'+1}, \dots, \mathcal{A}_{r-1}$ and the non-strict leaves corresponding to $\mathcal{L}_r, \mathcal{L}_{r+1}, \dots, \mathcal{L}_k$; call this intersection \mathcal{O}'' . It is clear that $\mathcal{O} \subseteq \mathcal{O}'' \subseteq \mathcal{O}'$ (since each axis is a subset of its corresponding non-strict leaf) and hence $\mathcal{O}'' \neq \{\mathbf{o}\}$, and so from Lemma 6.2.4 we see that \mathcal{O}'' is an original $(r'', k - r'')$ -piece for some $r'' \geq r - 1$.

However, Lemma 6.2.3 shows the dimension of \mathcal{O} to be $d - 2r - (k - r) = d - k - r$, and similarly the dimension of \mathcal{O}'' is $d - k - r''$. Corollary 6.2.7 shows \mathcal{O} and \mathcal{O}'' to be distinct and so Corollary 6.2.6 shows $\mathcal{O} = \mathcal{O} \cap \mathcal{O}''$ to have dimension strictly less than that of \mathcal{O}'' . So $r'' < r$ and we must have $r'' = r - 1$, and hence we have found an original $(r - 1, k - (r - 1))$ -piece \mathcal{O}'' containing \mathcal{O} .

Thus the maximal faces are in fact precisely the original $(r, k - r)$ -pieces \mathcal{O} for which $\mathcal{O} \not\subseteq \mathcal{O}'$ for any original $(r - 1, k - (r - 1))$ -piece \mathcal{O}' , noting that this condition is trivially satisfied if $r = 0$.

In conclusion consider the case $k = d - 1$. From Lemma 6.2.3 an original $(r, k - r)$ -piece must have $2r + (k - r) < d$ and so $r < d - k = 1$. So the only original $(r, k - r)$ -pieces are the original $(0, k)$ -pieces, all of which trivially satisfy the above condition. \blacksquare

Corollary 6.2.9 Consider a collection of k balanced $(d - 1)$ -fans in general position in a convex d -dimensional space \mathcal{P} , all of whose axes pass through some point $\mathbf{o} \in \text{int}(\mathcal{P})$, where $1 \leq k \leq d - 1$. Then if $\Phi > 1$, the maximal face count Φ is bounded above by

$$\Phi \leq \sum_{r=0}^{r_0} G_{r, k-r} \quad (6.2)$$

where $r_0 = \min(k, d - k - 1)$. Furthermore if k is as large as possible, i.e., $k = d - 1$, then Φ is given precisely by $\Phi = G_{0, k}$.

Proof From Theorem 6.2.8 we see that each maximal face is an original $(r, k - r)$ -piece for some r ; note that this requires $0 \leq r \leq k$. Furthermore from Lemma 6.2.3 we see that each original $(r, k - r)$ -piece must have $2r + (k - r) < d$, i.e., $r < d - k$. Thus each maximal face is an original $(r, k - r)$ -piece for $0 \leq r \leq \min(k, d - k - 1)$ and Equation 6.2 follows.

In the special case $k = d - 1$ we see from Theorem 6.2.8 that the maximal faces are precisely the original $(0, k)$ -pieces and so $\Phi = G_{0,k}$. ■

Note that for $k \leq d - 2$ the bound described by Equation 6.2 is quite loose since each original $(r, k - r)$ -piece will have a variety of original $(r', k - r')$ -pieces as subfaces, as described in Lemma 6.2.5. Thus most original $(r, k - r)$ -pieces will in fact not be maximal faces.

Nevertheless this bound may still be of use. Further discussion of this issue can be found in Section 6.6.2 at the end of this chapter.

6.3 Piece Count Calculations

It follows from Theorem 6.2.8 and Corollary 6.2.9 that our task must be to calculate or bound the piece counts $G_{r,s}$ for which $r + s = k$. For some piece counts it is straightforward to calculate an exact value or at least provide upper and lower bounds, as seen in the following set of results.

Lemma 6.3.1 Consider again the scenario described in Definitions 6.2.1.

1. If $r + s > k$ or $2r + s \geq d$ then $G_{r,s} = F_{r,s} = 0$.
2. If $r + s = k$ then $G_{r,s} = F_{r,s}$.
3. If $r \leq k$ and $2r < d$ then $G_{r,0} = \binom{k}{r}$.
4. If $s \geq 1$, $r + s \leq k$ and $2r + 2s \leq d$ then $G_{r,s} = \binom{k}{r} \binom{k-r}{s} 3^s$.
5. If $r + s \leq k$ then $G_{r,s} \leq F_{r,s} \leq 3^{k-r-s} G_{r,s}$.
6. The region count satisfies $R \leq 3^k$.

Proof We examine each statement in turn.

1. Recall from Definitions 6.2.1 that each original (r, s) -piece is constructed from the leaves and axes of $r + s$ distinct fans; thus there can be no original (r, s) -pieces if $r + s > k$.

Furthermore, if $s > 0$ and $2r + s \geq d$ then general position implies that s strict leaves and r axes from $r + s$ distinct fans cannot have non-empty intersection. If $s = 0$ and $2r + s \geq d$ then $2r \geq d$ and so general position implies that r axes from distinct fans can only intersect in the single point \mathbf{o} . Either way we again have no original (r, s) -pieces.

So if $r + s > k$ or $2r + s \geq d$ then $G_{r,s} = 0$, and since final (r, s) -pieces are all portions of original (r, s) -pieces it follows that $F_{r,s} = 0$.

2. Recall that each final (r, s) -piece is created by the subdivision of some original (r, s) -piece \mathcal{O} by the remaining $k - r - s$ fans not used in forming \mathcal{O} . If $r + s = k$ then there are no remaining fans and so \mathcal{O} is not subdivided at all.

Thus each original (r, s) -piece \mathcal{O} corresponds to just one final (r, s) -piece which is in fact all of \mathcal{O} , and we see that $G_{r,s} = F_{r,s}$.

3. If $2r < d$, general position implies that any r axes from distinct fans intersect in a $(d - 2r)$ -dimensional hyperplane. Thus, since $d - 2r \geq 1$, every choice of r distinct fans gives rise to an original $(r, 0)$ -piece. If $r \leq k$ then we can choose r distinct fans in $\binom{k}{r}$ different ways and so $G_{r,0} = \binom{k}{r}$.
4. We claim that if $s \geq 1$, $r + s \leq k$ and $2r + 2s \leq d$ then every possible selection of r axes and s leaves from $r + s$ distinct fans forms an original (r, s) -piece. This will be proven first for $s = 1$ and then for $s > 1$.

Assume that $s = 1$ and consider some selection of r axes $\mathcal{A}_1, \dots, \mathcal{A}_r$ and one strict leaf \mathcal{L} from $r + 1$ distinct fans. Since $2r + 2s \leq d$ we see that $2r < d$ and so by general position our r axes $\mathcal{A}_1, \dots, \mathcal{A}_r$ intersect in some $(d - 2r)$ -dimensional hyperplane \mathcal{H}_A . If \mathcal{H}_A meets \mathcal{L} then together they form the original (r, s) -piece that we require.

Suppose then that \mathcal{H}_A does not meet \mathcal{L} . Let \mathcal{H}_L denote the $(d - 1)$ -dimensional hyperplane containing \mathcal{L} and let \mathcal{A}_L denote the axis corresponding to leaf \mathcal{L} . We know that hyperplanes \mathcal{H}_A , \mathcal{H}_L and \mathcal{A}_L must intersect (they all contain central point \mathbf{o}). Since \mathcal{A}_L divides \mathcal{H}_L into two half-spaces one of which represents the leaf \mathcal{L} , it follows that the only way in which \mathcal{H}_A can avoid meeting strict leaf \mathcal{L} is if $\mathcal{H}_A \subseteq \mathcal{A}_L$.

In this case however, the intersection of all $r + 1$ axes $\mathcal{A}_1, \dots, \mathcal{A}_r, \mathcal{A}_L$ is also the $(d - 2r)$ -dimensional space \mathcal{H}_A , contradicting general position. So \mathcal{H}_A does indeed meet \mathcal{L} and we have our original (r, s) -piece as described above.

Assume now that $s > 1$ and consider some selection of r axes and s leaves from $r + s$ distinct fans. Without loss of generality let the s strict leaves be $\mathcal{L}_1, \dots, \mathcal{L}_s$ from fans $\mathcal{F}_1, \dots, \mathcal{F}_s$ respectively and let the r axes be from fans $\mathcal{F}_{s+1}, \dots, \mathcal{F}_{r+s}$. Let \mathcal{A}_i denote the axis of fan \mathcal{F}_i for each i .

Then for each j in the range $1 \leq j \leq s$ we know from the previous case $s = 1$ that the leaf \mathcal{L}_j and the remaining axes $\mathcal{A}_1, \dots, \mathcal{A}_{j-1}, \mathcal{A}_{j+1}, \dots, \mathcal{A}_{r+s}$ (i.e., all axes excluding \mathcal{A}_j) together form an original $(r + s - 1, 1)$ -piece. In particular there is some point \mathbf{p}_j belonging to strict leaf \mathcal{L}_j and each axis \mathcal{A}_i for $i \neq j$.

Consider now the centroid $\mathbf{p} = \frac{1}{n}(\mathbf{p}_1 + \dots + \mathbf{p}_{r+s})$. Since each \mathbf{p}_i belongs to each axis $\mathcal{A}_{s+1}, \dots, \mathcal{A}_{r+s}$ it follows that \mathbf{p} also belongs to each of these axes. Furthermore, for each j in the range $1 \leq j \leq s$ we see that \mathbf{p}_j belongs to strict leaf \mathcal{L}_j and each other \mathbf{p}_i for $i \neq j$ belongs to axis \mathcal{A}_j . Thus \mathbf{p} also belongs to the strict leaf \mathcal{L}_j .

So \mathbf{p} belongs to the intersection of our original s strict leaves and r axes (and $\mathbf{p} \neq \mathbf{o}$ since no strict leaf includes \mathbf{o}), and thus these leaves and axes form an original (r, s) -piece.

Thus our claim is proven for all s , and it follows that $G_{r,s}$ is simply the number of ways of choosing r axes and s leaves from $r + s$ distinct fans, which is $\binom{k}{r} \binom{k-r}{s} 3^s$.

5. It is clear that $G_{r,s} \leq F_{r,s}$ since the final (r, s) -pieces are subdivisions of the original (r, s) -pieces. Now consider these subdivisions in detail.

Each balanced fan subdivides the entire polytope into three convex regions. Since each original (r, s) -piece \mathcal{O} is also convex, it follows that each balanced fan not used in creating

\mathcal{O} will subdivide \mathcal{O} into at most three different convex portions. Applying this argument to each remaining fan in turn, we see that together the $k - r - s$ fans not used in creating \mathcal{O} will subdivide \mathcal{O} into $\leq 3^{k-r-s}$ convex portions.

Thus each original (r, s) -piece can give rise to at most 3^{k-r-s} final (r, s) -pieces, and so $F_{r,s} \leq 3^{k-r-s} G_{r,s}$.

6. Recalling from Definitions 6.2.1 that $R = F_{0,0}$, we can use the previous result with $r = s = 0$ to obtain $R = F_{0,0} \leq 3^k G_{0,0}$. Since the only original $(0, 0)$ piece is the entire polytope \mathcal{P} , we have $G_{0,0} = 1$ and so $R \leq 3^k$.

This concludes the series of proofs. ■

Recall from Theorem 6.2.8 and Corollary 6.2.9 that our ultimate aim is to calculate or bound the piece counts $G_{r,s}$ for which $r + s = k$. Whilst the equations and inequalities presented in Lemma 6.3.1 offer a place at which to start, they are by no means sufficient for calculating the piece counts that remain unknown. For this we will create families of equations derived from the Euler characteristics of various structures within our arrangement of fans.

Definitions 6.3.2 (Diagonal Sums) Consider a collection of k balanced $(d - 1)$ -fans in general position in a convex d -dimensional space \mathcal{P} , all of whose axes pass through some point $\mathbf{o} \in \text{int}(\mathcal{P})$, where $1 \leq k \leq d - 1$.

For each h satisfying $0 \leq h \leq k$ we define the *original h -diagonal sum* σ_h and the *final h -diagonal sum* τ_h as follows.

$$\begin{aligned}\sigma_h &= \sum_{\substack{r+s=h \\ 2r+s < d}} (-1)^s G_{r,s}; \\ \tau_h &= \sum_{\substack{r+s=h \\ 2r+s < d}} (-1)^s F_{r,s}.\end{aligned}$$

Example 6.3.3 Consider the piece counts calculated in Example 6.2.2 for which $d = 3$ and $k = 2$, summarised in Table 6.1. The corresponding diagonal sums are calculated as follows.

$$\begin{array}{llll} \sigma_0 & = & G_{0,0} & = & 1 & \quad \tau_0 & = & F_{0,0} & = & 9 \\ \sigma_1 & = & G_{1,0} - G_{0,1} & = & -4 & \quad \tau_1 & = & F_{1,0} - F_{0,1} & = & -12 \\ \sigma_2 & = & G_{0,2} & = & 5 & \quad \tau_2 & = & F_{0,2} & = & 5 \end{array}$$

Lemma 6.3.4 The following facts are true of the diagonal sums introduced in Definitions 6.3.2.

1. $\tau_0 = R$;
2. $\tau_k = \sigma_k$;
3. If $\Phi > 1$ and $k = d - 1$ then $\Phi = (-1)^k \tau_k = (-1)^k \sigma_k$.

Proof We prove each claim in turn.

1. From Definitions 6.3.2 we see that $\tau_0 = F_{0,0}$, and from Definitions 6.2.1 we see that $F_{0,0} = R$. Thus $\tau_0 = R$.
2. Lemma 6.3.1 tells us that $F_{r,s} = G_{r,s}$ when $r + s = k$. Since the definitions of τ_k and σ_k are identical except for the replacement of $F_{r,s}$ with $G_{r,s}$ where $r + s = k$, it follows that $\tau_k = \sigma_k$.
3. Let $k = d - 1$. If we have a pair (r, s) satisfying $r + s = k$ and $2r + s < d$, we can subtract these two equations to obtain $r < d - k$, i.e., $r < 1$. So the only such pair (r, s) is $(0, k)$, and we see from Definitions 6.3.2 that $\sigma_k = (-1)^k G_{0,k}$.

In addition we recall from Corollary 6.2.9 that $\Phi = G_{0,k}$ when $\Phi > 1$ and $k = d - 1$ and from an earlier result of this lemma that $\tau_k = \sigma_k$. Thus $\Phi = (-1)^k \tau_k = (-1)^k \sigma_k$.

This concludes the proof of our claims. ■

Lemma 6.3.5 (Euler Characteristic Equations) Consider the scenario described in Definitions 6.3.2, and consider any h satisfying $0 \leq h \leq k$. Then the following equation holds.

$$\sigma_h = \sum_{i=h}^k \binom{i}{h} \tau_i + \begin{cases} (-1)^d \binom{k}{h} & \text{if } 2h < d \leq 2k; \\ 0 & \text{otherwise.} \end{cases}$$

Proof Consider the entire d -dimensional space \mathcal{P} with its boundary identified to a single point (call this ∞), forming a topological d -sphere. This structure can be seen as a trivial cell complex with one d -dimensional face (the entire space \mathcal{P}) and one 0-dimensional face (the vertex ∞) and we can calculate its Euler characteristic χ ; here $\chi = (-1)^d + (-1)^0 = (-1)^d + 1$ which is 2 for even d and 0 for odd d as expected for a d -sphere. Call this cell complex \mathcal{D} .

We can subdivide this cell complex by inserting a single fan into \mathcal{P} . For any fan we let the *regions of the fan* refer to the closures of the three convex d -dimensional regions into which the fan divides \mathcal{P} ; thus each region of a fan includes two leaves and the axis of the fan. Inserting some fan \mathcal{F} into \mathcal{P} thus produces a new cell complex with three d -dimensional faces (the three regions of \mathcal{F}), three $(d - 1)$ -dimensional faces (the three non-strict leaves of \mathcal{F}), one $(d - 2)$ -dimensional face (the axis of \mathcal{F}) and one 0-dimensional face (the vertex ∞), again producing Euler characteristic $\chi = 3 \cdot (-1)^d + 3 \cdot (-1)^{d-1} + (-1)^{d-2} + (-1)^0 = (-1)^d + 1$ as expected. Call this new subcomplex $\mathcal{D}(\mathcal{F})$.

In general consider the cell complex $\mathcal{D}(\mathcal{F}_1, \dots, \mathcal{F}_m)$ formed by inserting fans $\mathcal{F}_1, \dots, \mathcal{F}_m$ into \mathcal{P} . This cell complex is the simplest common subdivision of $\mathcal{D}(\mathcal{F}_1), \dots, \mathcal{D}(\mathcal{F}_m)$ and so the faces of $\mathcal{D}(\mathcal{F}_1, \dots, \mathcal{F}_m)$ are precisely the intersections of one face from each $\mathcal{D}(\mathcal{F}_i)$. That is, the faces of $\mathcal{D}(\mathcal{F}_1, \dots, \mathcal{F}_m)$ are precisely

- the single vertex ∞ ;
- each intersection of r axes, s non-strict leaves and t regions from all $r + s + t$ distinct fans in our collection where $r + s + t = m$.

Let our collection of k fans be $\mathcal{F}_1, \dots, \mathcal{F}_k$ and consider some subset S containing h of our k fans; without loss of generality let $S = \{\mathcal{F}_1, \dots, \mathcal{F}_h\}$. We will form two new types of cell complex and calculate their Euler characteristics as follows.

First consider the cell complex $\mathcal{D}(\mathcal{F}_1, \dots, \mathcal{F}_h)$. This complex can be restricted to include only the points belonging to all h fans; this produces a new cell complex \mathcal{H}_S . It follows that the faces of \mathcal{H}_S are precisely

- the single vertex ∞ ;
- each intersection of r axes and s non-strict leaves from all $r + s$ distinct fans of S where $r + s = h$.

From Lemma 6.2.4 it follows that the faces of \mathcal{H}_S are precisely the vertex ∞ , the original (r, s) -pieces formed from the fans in S where $r + s = h$ and possibly the vertex \mathbf{o} , and from Corollary 6.2.7 we see that these faces are all distinct.

In particular, if some intersection of r axes and s leaves is the single point \mathbf{o} then the corresponding intersection of $r + s$ axes is also the single point \mathbf{o} , so vertex \mathbf{o} is a face of \mathcal{H}_S if and only if the intersection of all h axes in S is simply \mathbf{o} , i.e., if and only if $2h \geq d$ (using general position).

So consider the Euler characteristic $\chi(\mathcal{H}_S)$ and sum these Euler characteristics over all sets S containing h fans. Each original (r, s) -piece where $r + s = h$ will appear in precisely one of these subcomplexes, and the vertex ∞ and the vertex \mathbf{o} (if it is present) will appear in all $\binom{k}{h}$ of these complexes. Thus the sum of Euler characteristics is

$$\sum_S \chi(\mathcal{H}_S) = \binom{k}{h} + \sum_{\substack{r+s=h \\ 2r+s < d}} (-1)^{d-2r-s} G_{r,s} + \begin{cases} \binom{k}{h} & \text{if } 2h \geq d; \\ 0 & \text{otherwise,} \end{cases} \quad (6.3)$$

using Lemma 6.2.3 to calculate the dimension of each original (r, s) -piece.

As our second task consider the cell complex $\mathcal{D}(\mathcal{F}_1, \dots, \mathcal{F}_k)$ formed by inserting all k fans into \mathcal{P} . This can again be restricted to include only the points belonging to the h fans in S ; call this new cell complex \mathcal{K}_S . The faces of \mathcal{K}_S are then

- the single vertex ∞ ;
- each intersection of r axes, s non-strict leaves and t regions from all $r + s + t$ distinct fans where $r + s + t = k$ and the t regions all come from the $k - h$ fans not in S .

Consider some face \mathcal{I} of the second type described above. For each of the t regions used, if \mathcal{I} contains no points in the region interior then we can replace this region with a non-strict leaf bordering the region (since \mathcal{I} is convex we do not need both bordering leaves), and so \mathcal{I} can be written as an intersection of r axes, s non-strict leaves and t regions from $r + s + t$ distinct fans where $r + s + t = k$, the t regions all come from fans not in S and \mathcal{I} contains a point from each of the corresponding t region interiors.

Let the intersection of the r axes and s non-strict leaves be \mathcal{O} ; from Lemma 6.2.4 we see that \mathcal{O} is an original (r', s') -piece where $r' + s' = r + s$ (or is just \mathbf{o} , but we will ignore this case for now). Furthermore, using an argument similar to the proof of Lemma 6.1.3 we can convert our t points of \mathcal{I} in the t region interiors into a single point of \mathcal{I} that simultaneously belongs to all t region interiors (this new point being the centroid of the original t points). Thus the intersection of all t region interiors is non-empty and hence is an open d -dimensional space; call this \mathcal{R} . Since \mathcal{O} meets \mathcal{R} it follows that $\mathcal{O} \cap \mathcal{R}$ and hence $\mathcal{I} = \mathcal{O} \cap \mathcal{R}$ is a space of the same dimension as \mathcal{O} .

Thus we see that \mathcal{I} is a final (r', s') -piece formed from the subdivision of original piece \mathcal{O} by the remaining t fans. Recall that the set of fans providing the r' axes and s' leaves used in forming final piece \mathcal{I} is a superset of S .

Conversely it is easy to see that each final (r, s) -piece whose r axes and s leaves form a superset of S is in fact an intersection of r axes, s non-strict leaves and $t = k - r - s$ regions from all k distinct fans whose regions all come from the $k - h$ fans not in S , and is thus a face of \mathcal{K}_S .

It follows that the faces of \mathcal{K}_S are precisely ∞ , the final pieces satisfying the aforementioned property and possibly the vertex \mathbf{o} . From Corollary 6.2.7 these faces are again distinct, and by a similar argument as was used earlier, vertex \mathbf{o} appears if and only if all k axes intersect in the single point \mathbf{o} , i.e., if and only if $2k \geq d$.

Again we can sum the Euler characteristic $\chi(\mathcal{K}_S)$ over all sets S containing h fans. This time each final (r, s) -piece will be counted precisely $\binom{r+s}{h}$ times (the number of ways of choosing h fans that are a subset of the $r + s$ fans providing the axes and leaves) and so we obtain

$$\sum_S \chi(\mathcal{K}_S) = \binom{k}{h} + \sum_{\substack{r+s \geq h \\ 2r+s < d}} (-1)^{d-2r-s} \binom{r+s}{h} F_{r,s} + \begin{cases} \binom{k}{h} & \text{if } 2k \geq d; \\ 0 & \text{otherwise.} \end{cases} \quad (6.4)$$

At this point we observe that for each S , cell complexes \mathcal{H}_S and \mathcal{K}_S each represent the same underlying space (the intersection of the h fans in S with the boundary of \mathcal{P} identified to ∞), and so their Euler characteristics are equal. Thus the sums in Equations 6.3 and 6.4 must equate and so

$$\sum_{\substack{r+s=h \\ 2r+s < d}} (-1)^{d-2r-s} G_{r,s} = \sum_{\substack{r+s \geq h \\ 2r+s < d}} (-1)^{d-2r-s} \binom{r+s}{h} F_{r,s} + \begin{cases} \binom{k}{h} & \text{if } 2h < d \leq 2k; \\ 0 & \text{otherwise.} \end{cases}$$

Thus

$$(-1)^d \sigma_h = \sum_{i=h}^k (-1)^d \tau_i + \begin{cases} \binom{k}{h} & \text{if } 2h < d \leq 2k; \\ 0 & \text{otherwise,} \end{cases}$$

leading to the desired result. ■

Example 6.3.6 Consider the case where $d = 3$ and $k = 2$. We will aim to find an upper bound for the maximal face count Φ . Assume then that $\Phi > 1$.

From Lemma 6.3.4 we have $\Phi = \sigma_2$ and so we try applying Lemma 6.3.5, leading to the following results.

$$\sigma_0 = \tau_0 + \tau_1 + \tau_2 - \binom{2}{0}; \quad (6.5)$$

$$\sigma_1 = \tau_1 + 2\tau_2 - \binom{2}{1}; \quad (6.6)$$

$$\sigma_2 = \tau_2. \quad (6.7)$$

Recall from Lemma 6.3.1 that we know explicit values for far more of the original piece counts $G_{r,s}$ than of the final piece counts $F_{r,s}$. Thus it makes sense to use the above equations to eliminate as many of the final diagonal sums τ_i as possible.

From Equation 6.7 we have $\tau_2 = \sigma_2 = \Phi$, and so with Equation 6.6 we have $\tau_1 = \sigma_1 - 2\Phi + 2$. From Equation 6.5 then we have $\tau_0 = \sigma_0 - \tau_1 - \Phi + 1 = \sigma_0 - \sigma_1 + \Phi - 1$.

Recall however from Lemma 6.3.4 that $\tau_0 = R$, our region count. Thus we have an explicit relationship between R and Φ . Using Lemma 6.3.1 we can fill in values for σ_0 and σ_1 as follows.

$$\begin{aligned}
R &= \sigma_0 - \sigma_1 + \Phi - 1 \\
&= G_{0,0} - G_{1,0} + G_{0,1} + \Phi - 1 \\
&= 1 - 2 + 6 + \Phi - 1 \\
&= \Phi + 4.
\end{aligned}$$

So $R = \Phi + 4$. But from Lemma 6.3.1 we know that $R \leq 3^k = 9$, and so $\Phi \leq 9 - 4 = 5$. Thus we have shown that our number of maximal faces is $\leq 5 = 2^{d-1} + 1$, precisely the bound suggested in Conjecture 5.5.14.

We now wish to generalise the procedure used in Example 6.3.6 to obtain a relationship between region counts and maximal face counts for all d . We will begin as before by taking the equations of Lemma 6.3.5 and eliminating τ_i for $i > 0$. Recall that when $i = 0$ we have $\tau_0 = R$, our region count.

Before this however we prove a small result that will assist with the algebra that lies ahead.

Lemma 6.3.7 Consider some set of $2(m+1)$ values x_0, \dots, x_m and y_0, \dots, y_m . If it is true that

$$x_t = \sum_{i=t}^m \binom{i}{t} y_i \quad (6.8)$$

for each t satisfying $0 \leq t \leq m$ then it is also true that

$$y_u = \sum_{i=u}^m (-1)^{i-u} \binom{i}{u} x_i \quad (6.9)$$

for each u satisfying $0 \leq u \leq m$.

Proof We prove this by reverse induction on u . If $u = m$ then from Equation 6.8 we have $x_m = \binom{m}{m} y_m = y_m$, and so $y_m = x_m = \binom{m}{m} x_m$.

Now for some $k < m$ assume Equation 6.9 is true for each $u > k$ and consider the case $u = k$. Since $\binom{k}{k} = 1$, Equation 6.8 gives us $x_k = y_k + \sum_{i=k+1}^m \binom{i}{k} y_i$ and so

$$\begin{aligned}
y_k &= x_k - \sum_{i=k+1}^m \binom{i}{k} y_i \\
&= x_k - \sum_{i=k+1}^m \sum_{j=i}^m (-1)^{j-i} \binom{j}{i} \binom{i}{k} x_j,
\end{aligned}$$

using Equation 6.9 for each $u > k$. Furthermore, note that

$$\binom{j}{i} \binom{i}{k} = \frac{j!}{(j-i)!(i-k)!k!} = \binom{j}{k} \binom{j-k}{j-i}$$

and so

$$\begin{aligned}
y_k &= x_k - \sum_{i=k+1}^m \sum_{j=i}^m (-1)^{j-i} \binom{j}{k} \binom{j-k}{j-i} x_j \\
&= x_k - \sum_{j=k+1}^m (-1)^j \binom{j}{k} x_j \sum_{i=k+1}^j (-1)^i \binom{j-k}{j-i} \\
&= x_k + (-1)^k \sum_{j=k+1}^m (-1)^j \binom{j}{k} x_j
\end{aligned}$$

using the fact that $\sum_{i=k}^j (-1)^i \binom{j-k}{j-i} = (-1)^j \sum_{i=0}^{j-k} (-1)^i \binom{j-k}{i} = 0$ by a well-known identity and so $\sum_{i=k+1}^j (-1)^i \binom{j-k}{j-i} = -(-1)^k \binom{j-k}{j-k} = -(-1)^k$. Hence

$$y_k = x_k + \sum_{j=k+1}^m (-1)^{j-k} \binom{j}{k} x_j = \sum_{j=k}^m (-1)^{j-k} \binom{j}{k} x_j,$$

establishing Equation 6.9 for $u = k$. By induction it follows that Equation 6.9 is true for all u in the range $0 \leq u \leq m$. ■

Lemma 6.3.8 In the scenario described in Definitions 6.3.2, define

$$\sigma'_h = \begin{cases} \sigma_h + (-1)^{d+1} \binom{k}{h} & \text{if } 2h < d \leq 2k; \\ \sigma_h & \text{otherwise} \end{cases}$$

for each h satisfying $0 \leq h \leq k$. Then

$$\tau_0 = R = \sum_{i=0}^k (-1)^i \sigma'_i.$$

Proof First we observe that the equations of Lemma 6.3.5 can be written as

$$\sigma'_h = \sum_{i=h}^k \binom{i}{h} \tau_i$$

for each h . We can thus apply Lemma 6.3.7 using $x_i = \sigma'_i$ and $y_i = \tau_i$ to obtain

$$\tau_0 = \sum_{i=0}^k (-1)^i \binom{i}{0} \sigma'_i = \sum_{i=0}^k (-1)^i \sigma'_i.$$

Our proof is completed by recalling from Lemma 6.3.4 that $\tau_0 = R$. ■

Example 6.3.9 Consider now the case where $d = 4$ and $k = 3$, and again assume our maximal face count Φ satisfies $\Phi > 1$. Recall from Lemma 6.3.4 that $\Phi = -\sigma_3$.

Once more we aim to relate Φ to the region count R . From Lemma 6.3.8 we have $R = (\sigma_0 - \binom{3}{0}) - (\sigma_1 - \binom{3}{1}) + \sigma_2 - \sigma_3$. Using Lemma 6.3.1 to fill in explicit values for σ_0 , σ_1 and σ_2

we obtain

$$\begin{aligned}
R &= (\sigma_0 - 1) - (\sigma_1 - 3) + \sigma_2 + \Phi \\
&= G_{0,0} - G_{1,0} + G_{0,1} - G_{1,1} + G_{0,2} + 2 + \Phi \\
&= 1 - 3 + 9 - 18 + 27 + 2 + \Phi \\
&= \Phi + 18.
\end{aligned}$$

From Lemma 6.3.1 we know that $R \leq 3^3 = 27$ and so $\Phi \leq 27 - 18 = 9$. Thus our maximal face count is bounded above by $9 = 2^{d-1} + 1$, again attaining the exact bound suggested in Conjecture 5.5.14.

Unfortunately the situation worsens for $d \geq 5$ since we can no longer fill in all of the required piece counts $G_{r,s}$ with exact values from Lemma 6.3.1. Instead we must rely upon combinations of Lemma 6.3.8 with itself in a variety of different situations, as described in the following section.

6.4 Partial Intersections

In the cases where $(d, k) = (3, 2)$ and $(4, 3)$, seen in Examples 6.3.6 and 6.3.9 respectively, we relied upon being able to explicitly calculate $G_{r,s}$ whenever $r + s < k$. This was only possible through a happy accident in which $2r + 2s \leq d$ whenever $r + s < k$, allowing us to use Lemma 6.3.1 to provide exact values for $G_{r,s}$.

Once we reach $d \geq 5$ however there are values $G_{r,s}$ that we require but that Lemma 6.3.1 cannot immediately compute. For instance, if $k = d - 1$ then we cannot compute $G_{0,k-1}$ for $d \geq 5$ since we will have $2r + 2s = 2(k - 1) = 2d - 4 \not\leq d$. So we see that the techniques used in Examples 6.3.6 and 6.3.9 must somehow be extended.

Our approach to this problem is to apply the results already obtained to intersections of only some of our k fans, thus obtaining new formulae that can help eliminate more of the unknown diagonal sums σ_i .

Definitions 6.4.1 (Partial Counts and Sums) Consider a collection of k balanced $(d-1)$ -fans in general position in a convex d -dimensional space \mathcal{P} , all of whose axes pass through some point $\mathbf{o} \in \text{int}(\mathcal{P})$, where $1 \leq k \leq d - 1$.

For any h satisfying $1 \leq h \leq k$, consider some subset of precisely h of our k fans. The intersection of these h fans in \mathcal{P} produces its own region count, original and final (r, s) -pieces and original and final diagonal sums. We thus define the following quantities.

- The h -partial region count $R^{(h)}$ is the region count for the intersection of precisely h of our k fans, summed over all possible choices of these h fans.
- The h -partial original and final (r, s) -piece counts $G_{r,s}^{(h)}$ and $F_{r,s}^{(h)}$ are the original and final (r, s) -piece counts for the intersection of precisely h of our k fans, again summed over all possible choices of these h fans.

- The h -partial original and final l -diagonal sums $\sigma_l^{(h)}$ and $\tau_l^{(h)}$ are the original and final l -diagonal sums for the intersection of precisely h of our k fans, once more summed over all possible choices of these h fans.

Furthermore, we extend the definition of $R^{(h)}$ to include $R^{(0)} = 1$, signifying the fact that when 0 fans are chosen the polytope \mathcal{P} remains a single undivided region.

Example 6.4.2 Consider once more Example 6.2.2 for which $d = 3$ and $k = 2$, illustrated in Figure 6.2. We examine here the different 1-partial counts and sums.

Consider the intersection of only one of these two fans. This single fan has three original $(0, 1)$ -pieces (the three leaves) and one original $(1, 0)$ -piece (the axis). Since we are looking at the intersection of just one fan, none of these pieces are subdivided further and so we also have three final $(0, 1)$ -pieces and one final $(1, 0)$ -piece. Furthermore there is a single original $(0, 0)$ -piece (the entire polytope) which is subdivided into three final $(0, 0)$ -pieces (i.e., three regions) by the fan.

Thus for the intersection of just one fan we have the following results.

$$\begin{array}{llll}
 G_{0,0} = 1 & F_{0,0} = 3 & \sigma_0 = G_{0,0} = 1 & \tau_0 = F_{0,0} = 3 \\
 G_{1,0} = 1 & F_{1,0} = 1 & \sigma_1 = G_{1,0} - G_{0,1} = -2 & \tau_1 = F_{1,0} - F_{0,1} = -2 \\
 G_{0,1} = 3 & F_{0,1} = 3 & & R = F_{0,0} = 3
 \end{array}$$

Now consider our original intersection of two fans. Each 1-partial count or sum is constructed by summing the above counts and sums over all possible choices of just one fan. There are two ways of choosing one fan and so each of the above numbers is simply doubled, giving the following results.

$$\begin{array}{llll}
 G_{0,0}^{(1)} = 2 & F_{0,0}^{(1)} = 6 & \sigma_0^{(1)} = 2 & \tau_0^{(1)} = 6 \\
 G_{1,0}^{(1)} = 2 & F_{1,0}^{(1)} = 2 & \sigma_1^{(1)} = -4 & \tau_1^{(1)} = -4 \\
 G_{0,1}^{(1)} = 6 & F_{0,1}^{(1)} = 6 & & R^{(1)} = 6
 \end{array}$$

Lemma 6.4.3 Consider the scenario described in Definitions 6.4.1. For each h in the range $1 \leq h \leq k$ we have the following identities.

1. $G_{r,s}^{(h)} = \binom{k-r-s}{k-h} G_{r,s}$;
2. $\sigma_l^{(h)} = \binom{k-l}{k-h} \sigma_l$.

Furthermore, it is true that $G_{r,s}^{(k)} = G_{r,s}$, $F_{r,s}^{(k)} = F_{r,s}$, $\sigma_l^{(k)} = \sigma_l$, $\tau_l^{(k)} = \tau_l$ and $R^{(k)} = R$.

Proof Consider each set of identities in turn.

1. If $r + s > h$, it is clear from Lemma 6.3.1 that $G_{r,s}^{(h)} = 0$. But it is also true if $r + s > h$ that $\binom{k-r-s}{k-h} = 0$ and so the required identity holds.

Otherwise suppose that $r + s \leq h$. Note first of all that each original (r, s) -piece in the intersection of just h fans is also an original (r, s) -piece in the overall intersection of all k fans.

Now consider some original (r, s) -piece \mathcal{O} in the overall intersection of k fans, and consider some intersection \mathcal{H} using only h of our original k fans. Then intersection \mathcal{H} includes original piece \mathcal{O} if and only if the h fans used to create \mathcal{H} include the $r + s$ fans used to create \mathcal{O} .

The number of ways that we can choose an intersection \mathcal{H} that includes \mathcal{O} is then $\binom{k-r-s}{h-r-s}$, since although $r + s$ of the fans in \mathcal{H} are predetermined, we are free to choose the other $h - r - s$ fans from the remaining $k - r - s$ fans available.

Thus when we sum over all intersections that use just h of our original k fans, we find that each original (r, s) -piece is counted precisely $\binom{k-r-s}{h-r-s} = \binom{k-r-s}{k-h}$ times. So $G_{r,s}^{(h)} = \binom{k-r-s}{k-h} G_{r,s}$.

2. Since each $\sigma_l^{(h)}$ is by definition a linear combination of $G_{r,s}^{(h)}$ terms for which $r + s = l$, it is clear from the previous result that $\sigma_l^{(h)} = \binom{k-l}{k-h} \sigma_l$.

Finally, the only way of choosing k of the original k fans to intersect is to simply use the original intersection of all k fans. Thus it is clear that $G_{r,s}^{(k)} = G_{r,s}$, $F_{r,s}^{(k)} = F_{r,s}$, $\sigma_l^{(k)} = \sigma_l$, $\tau_l^{(k)} = \tau_l$ and $R^{(k)} = R$. ■

Lemma 6.4.4 In the scenario described in Definitions 6.4.1, define

$$\sigma'_l{}^{(h)} = \sigma_l^{(h)} + \begin{cases} (-1)^{d+1} \binom{k}{h} \binom{h}{l} & \text{if } 2l < d \leq 2h; \\ 0 & \text{otherwise} \end{cases}$$

for each h and l satisfying $0 \leq l \leq h \leq k$. Then

$$R^{(h)} = \sum_{i=0}^h (-1)^i \sigma_i'^{(h)}.$$

Proof Consider some intersection of only h of our original k fans. To this intersection we can apply Lemma 6.3.8 to obtain

$$R = \sum_{i=0}^h (-1)^i \left(\sigma_i + \begin{cases} (-1)^{d+1} \binom{k}{h} \binom{h}{i} & \text{if } 2i < d \leq 2h; \\ 0 & \text{otherwise,} \end{cases} \right)$$

where region count R and diagonal sums σ_i refer to this specific intersection of only h fans.

We can now sum over all choices of h fans. The region count and diagonal sums become the overall h -partial region count and h -partial diagonal sums, and the constant term is multiplied by $\binom{k}{h}$ since there are $\binom{k}{h}$ ways of choosing h of our original k fans. The result is

$$R^{(h)} = \sum_{i=0}^h (-1)^i \left(\sigma_i^{(h)} + \begin{cases} (-1)^{d+1} \binom{k}{h} \binom{h}{i} & \text{if } 2i < d \leq 2h; \\ 0 & \text{otherwise,} \end{cases} \right)$$

and so $R^{(h)} = \sum_{i=0}^h (-1)^i \sigma_i'^{(h)}$ as required. ■

Recall from Lemma 6.3.4 that when k is as large as possible, i.e., $k = d - 1$, then in order to calculate the maximal face count Φ it suffices to know σ_k . Our next step then is to derive a formula for σ_k in terms of the partial region counts $R^{(h)}$.

Lemma 6.4.5 In the scenario described in Definitions 6.4.1, it is true that

$$\sigma_k = \sum_{i=0}^k (-1)^i R^{(i)} + \sum_{\substack{0 \leq l < h \leq k \\ 2l < d \leq 2h}} (-1)^{d+h+l} \binom{k}{h} \binom{h}{l},$$

where the right hand sum is over all h and l satisfying the given conditions.

Proof For each h satisfying $0 \leq h \leq k$, define

$$R'^{(h)} = R^{(h)} + \sum_{\substack{0 \leq l < h \\ 2l < d \leq 2h}} (-1)^{d+l} \binom{k}{h} \binom{h}{l},$$

where the above sum is over all l satisfying the summation conditions. Note that if $2h < d$ then the summation conditions can never be satisfied and so this expression will reduce to $R'^{(h)} = R^{(h)}$.

Then from Lemma 6.4.4 we see that $R'^{(h)} = \sum_{i=0}^h (-1)^i \sigma_i^{(h)}$. Note that we have replaced the condition $l \leq h$ (taken from Lemma 6.4.4) with the new condition $l < h$; this makes no difference since conditions $2l < d$ and $d \leq 2h$ together imply $l < h$ regardless. Applying Lemma 6.4.3 we now have

$$R'^{(h)} = \sum_{i=0}^h (-1)^i \binom{k-i}{k-h} \sigma_i. \quad (6.10)$$

For each i we define $x_i = R'^{(k-i)}$ and $y_i = (-1)^{k-i} \sigma_{k-i}$. Equation 6.10 then becomes

$$x_t = \sum_{i=t}^k \binom{i}{t} y_i.$$

But this allows us to invoke Lemma 6.3.7 from which we obtain $y_0 = \sum_{i=0}^k (-1)^i x_i$. Translating back in terms of $R'^{(h)}$ and σ_i we find

$$(-1)^k \sigma_k = \sum_{h=0}^k (-1)^{k-h} R'^{(h)}.$$

Thus

$$\begin{aligned} \sigma_k &= \sum_{h=0}^k (-1)^h R'^{(h)} \\ &= \sum_{h=0}^k (-1)^h R^{(h)} + \sum_{h=0}^k (-1)^h \sum_{\substack{0 \leq l < h \\ 2l < d \leq 2h}} (-1)^{d+l} \binom{k}{h} \binom{h}{l} \\ &= \sum_{h=0}^k (-1)^h R^{(h)} + \sum_{\substack{0 \leq l < h \leq k \\ 2l < d \leq 2h}} (-1)^{d+h+l} \binom{k}{h} \binom{h}{l}. \end{aligned}$$

■

Lemma 6.4.6 For any integers d and k satisfying $1 \leq d$ and $1 \leq k$, it is true that

$$\sum_{\substack{0 \leq l < h \leq k \\ 2l < d \leq 2h}} (-1)^{h+l} \binom{k}{h} \binom{h}{l} = \begin{cases} -1 & \text{if } 2k \geq d; \\ 0 & \text{if } 2k < d. \end{cases}$$

Proof We prove this by induction on k . To begin, consider the case where $k = 1$. If $2k < d$ then it is impossible to simultaneously satisfy both conditions $h \leq k$ and $d \leq 2h$, and so the given sum is zero as required. If $2k \geq d$ then the only pair (h, l) satisfying the summation conditions is $(1, 0)$ and so the sum is $(-1)^{1+0} \binom{1}{1} \binom{1}{0} = -1$ again as required.

Now consider some arbitrary $k \geq 2$ and assume this lemma holds for all smaller values of k . Let S denote the sum we seek to evaluate, i.e.,

$$S = \sum_{\substack{0 \leq l < h \leq k \\ 2l < d \leq 2h}} (-1)^{h+l} \binom{k}{h} \binom{h}{l}.$$

If $d = 1$ or $d = 2$ then the only pairs (h, l) satisfying the summation conditions are $(h, 0)$ for $h > 0$, and so $S = \sum_{h=1}^k (-1)^h \binom{k}{h} = -1$ (using the well-known result that $\sum_{h=0}^k (-1)^h \binom{k}{h} = 0$). Note also that in these cases we have $2k \geq d$ and so $S = -1$ is indeed the desired result.

Otherwise we have $k \geq 2$ and $d \geq 3$. Using the identity $\binom{n}{r} = \binom{n-1}{r} + \binom{n-1}{r-1}$ with the $\binom{n-1}{r}$ term absent if $r = n$ and the $\binom{n-1}{r-1}$ term absent if $r = 0$, the sum S can be rewritten as

$$S = \sum_{\substack{0 \leq l < h \leq k-1 \\ 2l < d \leq 2h}} (-1)^{h+l} \binom{k-1}{h} \binom{h}{l} + \sum_{\substack{0 \leq l < h \leq k \\ 2l < d \leq 2h}} (-1)^{h+l} \binom{k-1}{h-1} \binom{h-1}{l} + \sum_{\substack{1 \leq l < h \leq k \\ 2l < d \leq 2h}} (-1)^{h+l} \binom{k-1}{h-1} \binom{h-1}{l-1}.$$

If we define $k' = k - 1$, $h' = h - 1$ and $l' = l - 1$, this expression becomes

$$S = \sum_{\substack{0 \leq l < h \leq k' \\ 2l < d \leq 2h}} (-1)^{h+l} \binom{k'}{h} \binom{h}{l} - \sum_{\substack{0 \leq l \leq h' \leq k' \\ 2l < d \leq 2h}} (-1)^{h'+l} \binom{k'}{h'} \binom{h'}{l} + \sum_{\substack{0 \leq l' < h' \leq k' \\ 2l' < d-2 \leq 2h'}} (-1)^{h'+l'} \binom{k'}{h'} \binom{h'}{l'},$$

and by our inductive hypothesis (since $k' < k$) this reduces to

$$S = - \sum_{\substack{0 \leq l \leq h' \leq k' \\ 2l < d, d-2 \leq 2h'}} (-1)^{h'+l} \binom{k'}{h'} \binom{h'}{l} + \begin{cases} -2 & \text{if } d \leq 2k'; \\ -1 & \text{if } d-2 \leq 2k' < d; \\ 0 & \text{if } 2k' < d-2. \end{cases} \quad (6.11)$$

Consider now the sum on the right hand side of Equation 6.11. We can split this sum into two partial sums by partitioning the pairs (h, l) that satisfy the summation constraints $0 \leq l \leq h' \leq k'$, $2l < d$ and $d-2 \leq 2h'$ into the following two categories.

- $0 \leq l < h' \leq k'$ and $2l < d \leq 2h'$. Here our inductive hypothesis simply gives the corresponding partial sum as

$$\sum_{\substack{0 \leq l < h' \leq k' \\ 2l < d \leq 2h'}} (-1)^{h'+l} \binom{k'}{h'} \binom{h'}{l} = \begin{cases} -1 & \text{if } d \leq 2k'; \\ 0 & \text{if } 2k' < d. \end{cases} \quad (6.12)$$

- $0 \leq l \leq h' \leq k'$, $2l < d$ and $d - 2 \leq 2h' < d$. If $2k' < d - 2$ then there are no pairs (h', l) satisfying these constraints; otherwise these constraints are satisfied for precisely the pairs (h', l) for which $0 \leq l \leq h' = \lceil \frac{d-2}{2} \rceil$ (where $\lceil x \rceil$ represents x rounded up to the nearest integer). So defining $d_0 = \lceil \frac{d-2}{2} \rceil$, we observe that when $2k' \not< d - 2$ our partial sum becomes

$$\sum_{l=0}^{d_0} (-1)^{d_0+l} \binom{k'}{d_0} \binom{d_0}{l} = 0,$$

again using the well-known identity $\sum_{l=0}^{d_0} (-1)^l \binom{d_0}{l} = 0$. Thus whether $2k' < d - 2$ or $d - 2 \leq 2k'$, our partial sum is always

$$\sum_{\substack{0 \leq l < h' \leq k' \\ 2l < d-2 \leq 2h' < d}} (-1)^{h'+l} \binom{k'}{h'} \binom{h'}{l} = 0. \quad (6.13)$$

We can now substitute Equations 6.12 and 6.13 into the sum on the right hand side of Equation 6.11, which reduces all of Equation 6.11 to

$$S = \begin{cases} -(-1) + -2 & \text{if } d \leq 2k'; \\ 0 + -1 & \text{if } d - 2 \leq 2k' < d; \\ 0 + 0 & \text{if } 2k' < d - 2. \end{cases}$$

Thus we see that our original sum is simply

$$\begin{aligned} \sum_{\substack{0 \leq l < h \leq k \\ 2l < d \leq 2h}} (-1)^{h+l} \binom{k}{h} \binom{h}{l} &= S = \begin{cases} -1 & \text{if } d - 2 \leq 2k'; \\ 0 & \text{if } 2k' < d - 2 \end{cases} \\ &= \begin{cases} -1 & \text{if } d \leq 2k; \\ 0 & \text{if } 2k < d, \end{cases} \end{aligned}$$

as required. ■

Corollary 6.4.7 Consider a collection of k balanced $(d - 1)$ -fans in general position in a convex d -dimensional space \mathcal{P} , all of whose axes pass through some point $\mathbf{o} \in \text{int}(\mathcal{P})$, where $1 \leq k \leq d - 1$. Then

$$\sigma_k = \sum_{i=0}^k (-1)^i R^{(i)} + \begin{cases} (-1)^{d+1} & \text{if } 2k \geq d; \\ 0 & \text{if } 2k < d. \end{cases} \quad (6.14)$$

Proof This is an immediate consequence of Lemmas 6.4.5 and 6.4.6. ■

Corollary 6.4.8 Consider again a collection of k balanced $(d - 1)$ -fans in a convex d -dimensional space \mathcal{P} as described above, where k is as large as possible, i.e., $k = d - 1$. If $\Phi > 1$ then the maximal face count Φ is given precisely by

$$\Phi = \sum_{h=0}^k (-1)^{k-h} R^{(h)} + 1.$$

Proof This is an immediate application of Lemma 6.3.4 to Corollary 6.4.7, observing that when $1 \leq k = d - 1$ it is always true that $2k \geq d$. ■

Thus we see that we have reduced our number of maximal faces to an alternating sum of partial region counts. This may seem like nothing more than an orgy of algebraic manipulation, but in fact we have brought ourselves to a point where we can start to bound components of this alternating sum that we cannot precisely calculate.

In particular, recall from Lemma 6.3.1 that the region count is bounded above by $R \leq 3^k$. We can extend this result to bound the partial region counts, and then attempt to piece these bounds together to form an overall bound for σ_k and hence the maximal face count Φ .

6.5 Region Counting Revisited

As described above, in order to use Corollary 6.4.7 to bound σ_k and hence the maximal face count, we must understand in more detail the behaviour of the partial region counts $R^{(i)}$.

Lemma 6.5.1 Consider a collection of k balanced $(d - 1)$ -fans in general position in a convex d -dimensional space \mathcal{P} , all of whose axes pass through some point $\mathbf{o} \in \text{int}(\mathcal{P})$, where $1 \leq k \leq d - 1$. Then for each h satisfying $0 \leq h \leq k$, the partial region count $R^{(h)}$ is bounded by

$$\binom{k}{h} \leq R^{(h)} \leq \binom{k}{h} 3^h.$$

Proof Consider any intersection of h of our k fans. For this particular intersection the region count R is bounded below by $1 \leq R$ since the surrounding polytope is divided into at least one region, and is bounded above by $R \leq 3^h$ as seen in Lemma 6.3.1.

For our overall intersection of k fans then, the partial region count $R^{(h)}$ is the sum of these smaller region counts over all possible choices for h of our k fans. Since there are $\binom{k}{h}$ such choices it follows that $\binom{k}{h} \leq R^{(h)} \leq \binom{k}{h} 3^h$. ■

Lemma 6.5.2 In the scenario described in Lemma 6.5.1, if $2h \leq d$ then $R^{(h)}$ in fact always attains its upper bound

$$R^{(h)} = \binom{k}{h} 3^h.$$

Proof To evaluate $R^{(h)}$, we can use Lemma 6.4.4 to expand $R^{(h)}$ as a combination of partial diagonal sums $\sigma_i^{(h)}$, Lemma 6.4.3 to replace the $\sigma_i^{(h)}$ with diagonal sums σ_i and then Definitions 6.3.2 to replace the σ_i with piece counts $G_{r,s}$. Finally, since $2h \leq d$ we will be able to use Lemma 6.3.1 to obtain explicit values for the terms $G_{r,s}$ and thus evaluate $R^{(h)}$.

We begin by splitting into cases $2h < d$ and $2h = d$.

- Suppose that $2h < d$. Then in Lemma 6.4.4 we see that $\sigma_i'^{(h)} = \sigma_i^{(h)}$ for all i and so

$R^{(h)} = \sum_{i=0}^h (-1)^i \sigma_i^{(h)}$. Thus performing our expansions as described above we obtain

$$\begin{aligned}
R^{(h)} &= \sum_{i=0}^h (-1)^i \sigma_i^{(h)} \\
&= \sum_{i=0}^h (-1)^i \binom{k-i}{k-h} \sigma_i \\
&= \sum_{i=0}^h \sum_{\substack{r+s=i \\ 2r+s < d}} (-1)^i \binom{k-i}{k-h} (-1)^s G_{r,s} \\
&= \sum_{i=0}^h \sum_{r+s=i} (-1)^{i+s} \binom{k-i}{k-h} \binom{k}{r} \binom{k-r}{s} 3^s,
\end{aligned} \tag{6.15}$$

noting that we can remove $2r + s < d$ from the summation conditions in Equation 6.15 since $r + s = i \leq h$ and $2h < d$ together imply $2r + s < d$.

- Otherwise we have $2h = d$. In this case Lemma 6.4.4 has $\sigma_i'^{(h)} \neq \sigma_i^{(h)}$ unless $i = h$, giving the following sequence of expansions.

$$\begin{aligned}
R^{(h)} &= \sum_{i=0}^h (-1)^i \sigma_i'^{(h)} \\
&= \sum_{i=0}^h (-1)^i \sigma_i^{(h)} + \sum_{i=0}^{h-1} (-1)^i (-1)^{d+1} \binom{k}{h} \binom{h}{i} \\
&= \sum_{i=0}^h (-1)^i \binom{k-i}{k-h} \sigma_i + \binom{k}{h} \sum_{i=0}^{h-1} (-1)^{d+i+1} \binom{h}{i} \\
&= \sum_{i=0}^h \sum_{\substack{r+s=i \\ 2r+s < d}} (-1)^i \binom{k-i}{k-h} (-1)^s G_{r,s} + \binom{k}{h} (-1)^{d+h},
\end{aligned} \tag{6.16}$$

using the fact that $\sum_{i=0}^h (-1)^i \binom{h}{i} = 0$. Examining the summation conditions in Equation 6.16 we see that since $r + s = i \leq h$ and $2h = d$, we always have $2r + s < d$ unless $i = h$ and $(r, s) = (h, 0)$. Thus we can drop the condition $2r + s < d$ and subtract away this additional term that we gain, giving

$$\begin{aligned}
R^{(h)} &= \sum_{i=0}^h \sum_{r+s=i} (-1)^{i+s} \binom{k-i}{k-h} \binom{k}{r} \binom{k-r}{s} 3^s - (-1)^h \binom{k}{h} + \binom{k}{h} (-1)^{d+h} \\
&= \sum_{i=0}^h \sum_{r+s=i} (-1)^{i+s} \binom{k-i}{k-h} \binom{k}{r} \binom{k-r}{s} 3^s
\end{aligned}$$

since $d = 2h$ is even.

Thus in both cases we obtain

$$\begin{aligned}
R^{(h)} &= \sum_{i=0}^h \sum_{r+s=i} (-1)^{i+s} \binom{k-i}{k-h} \binom{k}{r} \binom{k-r}{s} 3^s \\
&= \sum_{r+s \leq h} (-1)^r \binom{k}{r} \binom{k-r}{s} \binom{k-r-s}{k-h} 3^s \\
&= \sum_{r+s \leq h} (-1)^r \binom{k}{k-h} \binom{h}{r} \binom{h-r}{s} 3^s,
\end{aligned}$$

since both $\binom{k}{r} \binom{k-r}{s} \binom{k-r-s}{k-h}$ and $\binom{k}{k-h} \binom{h}{r} \binom{h-r}{s}$ are the number of ways of dividing k objects into four groups of size $r, s, h-r-s$ and $k-h$. Finally this becomes

$$\begin{aligned}
R^{(h)} &= \binom{k}{k-h} \sum_{r+s \leq h} \binom{h}{r} \binom{h-r}{s} (-1)^r 3^s 1^{h-r-s} \\
&= \binom{k}{k-h} (-1 + 3 + 1)^h
\end{aligned}$$

using the multinomial expansion of $(x + y + z)^h$ and so $R^{(h)} = \binom{k}{k-h} 3^h = \binom{k}{h} 3^h$. ■

Example 6.5.3 Consider again Example 6.3.9 in which we prove that the number of maximal faces in the overall intersection of fans is ≤ 9 when $(d, k) = (4, 3)$. We now prove this in a more direct fashion using partial region counts.

From Corollary 6.4.8 we see that the number of maximal faces Φ is $\left(\sum_{i=0}^3 (-1)^{3-i} R^{(i)}\right) + 1$. Using Lemma 6.5.1 to evaluate $R^{(h)}$ for $h \leq 2$ and Lemma 6.5.2 to bound $R^{(3)}$ this gives

$$\begin{aligned}
\Phi &= -R^{(0)} + R^{(1)} - R^{(2)} + R^{(3)} + 1 \\
&= -1 + 9 - 27 + R^{(3)} + 1 \\
&\leq -1 + 9 - 27 + 27 + 1 \\
&= 9.
\end{aligned}$$

Observe that $R^{(h)}$ attains its upper bound of $\binom{k}{h} 3^h$ precisely for small h and that in Example 6.5.3 the maximal number of faces is achieved precisely when the remaining partial region counts also attain their upper bounds. It therefore makes sense to rephrase Corollary 6.4.7 so that instead of measuring region counts, we find some way of measuring deviance from this maximal case.

Definitions 6.5.4 (Region Space) Consider a collection of k balanced $(d-1)$ -fans labelled $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k$ in general position in a convex d -dimensional space \mathcal{P} , all of whose axes pass through some point $\mathbf{o} \in \text{int}(\mathcal{P})$, where $1 \leq k \leq d-1$.

Observe that each fan \mathcal{F}_i divides the space \mathcal{P} into three convex regions; arbitrarily label these regions $\rho_{i,0}, \rho_{i,1}$ and $\rho_{i,2}$. Each $\rho_{i,j}$ is defined to be an open region, so the three regions $\rho_{i,0}, \rho_{i,1}$ and $\rho_{i,2}$ are pairwise disjoint and their union is $\mathcal{P} \setminus \mathcal{F}_i$.

Furthermore, observe that the union of all k fans divides the space \mathcal{P} into R regions (where R is the region count), each of which can be expressed as $\rho_{1,i_1} \cap \rho_{2,i_2} \cap \dots \cap \rho_{k,i_k}$ for some unique vector $(i_1, i_2, \dots, i_k) \in \mathbb{Z}_3^k$.

We now make the following definitions.

- The *region space* for fans $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k$, denoted by $\mathfrak{R}(\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k)$, is simply the space \mathbb{Z}_3^k . Each vector of \mathbb{Z}_3^k has a specific meaning as seen below.

If the order in which these fans are presented is clear from context (or is unimportant) then the region space may be abbreviated as $\mathfrak{R}(S)$ where S is the set of fans used. If the fans themselves are also clear from context then the region space may be abbreviated further as simply \mathfrak{R} .

- The *associated regions* of each vector $(i_1, i_2, \dots, i_k) \in \mathfrak{R}$ are the regions $\rho_{1,i_1}, \rho_{2,i_2}, \dots, \rho_{k,i_k}$.
- A *visible region* is a vector $(i_1, i_2, \dots, i_k) \in \mathfrak{R}$ whose intersection of associated regions is one of the R regions into which space \mathcal{P} is divided, i.e., for which $\rho_{1,i_1} \cap \rho_{2,i_2} \cap \dots \cap \rho_{k,i_k} \neq \emptyset$.
- A *missing region* is a vector $(i_1, i_2, \dots, i_k) \in \mathfrak{R}$ that is not a visible region, i.e., for which $\rho_{1,i_1} \cap \rho_{2,i_2} \cap \dots \cap \rho_{k,i_k} = \emptyset$.

For convenience we can extend this definition for the case $k = 0$, where $\mathfrak{R}(\emptyset) = \mathbb{Z}_3^0 = \{()\}$. This single empty vector is considered a visible region, reflecting the fact that when no fans are present then the entire space \mathcal{P} forms a region of itself.

Definitions 6.5.5 (Missing Region Counts) Consider again the scenario described in Definitions 6.5.4. The *missing region count*, denoted by M , is the number of missing regions in the region space \mathfrak{R} .

Furthermore, for any h satisfying $1 \leq h \leq k$, the *h -partial missing region count*, denoted by $M^{(h)}$, is the missing region count for each intersection of h of our k fans, summed over all possible choices of h of our k fans. That is, $M^{(h)}$ is the number of missing regions in the region space $\mathfrak{R}(H)$, summed over all possible subsets $H \subseteq \{\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k\}$ for which $|H| = h$.

Finally we define $M^{(0)} = 0$ to reflect the fact that when no fans are selected, the entire space \mathcal{P} is indeed a region of \mathcal{P} and is not missing. Note that this is consistent with the earlier definition of $\mathfrak{R}(\emptyset)$.

Lemma 6.5.6 In the scenario described in Definitions 6.5.5, it is true that $M = 3^k - R$. Furthermore, $M^{(h)} = \binom{k}{h} 3^h - R^{(h)}$ for each h satisfying $0 \leq h \leq k$.

Proof From Definitions 6.5.4 we see that the visible regions in the region space \mathfrak{R} correspond precisely to the R regions into which space \mathcal{P} is divided, and that the remaining vectors in \mathfrak{R} are precisely our missing regions. Thus $M = |\mathfrak{R}| - R = |\mathbb{Z}_3^k| - R = 3^k - R$.

Now consider some intersection of only h of our k fans. For this partial intersection the above result gives $M = 3^h - R$. Summing over all selections of h of our k fans, the counts M and R become h -partial counts $M^{(h)}$ and $R^{(h)}$, and the constant term is multiplied by $\binom{k}{h}$ (the number of such selections). Thus in our overall intersection of k fans we have $M^{(h)} = \binom{k}{h} 3^h - R^{(h)}$. ■

Corollary 6.5.7 In the scenario described in Definitions 6.5.5, $M^{(h)}$ is bounded by

$$0 \leq M^{(h)} \leq \binom{k}{h} (3^h - 1)$$

for each h satisfying $0 \leq h \leq k$. Furthermore, if $2h \leq d$ then we in fact have $M^{(h)} = 0$.

Proof This is an immediate application of Lemma 6.5.6 to Lemmas 6.5.1 and 6.5.2. ■

Corollary 6.5.8 Consider a collection of k balanced $(d-1)$ -fans in general position in a convex d -dimensional space \mathcal{P} , all of whose axes pass through some point $\mathbf{o} \in \text{int}(\mathcal{P})$, where $1 \leq k \leq d-1$. Then

$$\sigma_k = (-2)^k - \sum_{i=0}^k (-1)^i M^{(i)} + \begin{cases} (-1)^{d+1} & \text{if } 2k \geq d; \\ 0 & \text{if } 2k < d. \end{cases}$$

Proof From Corollary 6.4.7 we have

$$\sigma_k = \sum_{i=0}^k (-1)^i R^{(i)} + \begin{cases} (-1)^{d+1} & \text{if } 2k \geq d; \\ 0 & \text{if } 2k < d. \end{cases}$$

Applying Lemma 6.5.6 however we observe that

$$\begin{aligned} \sum_{i=0}^k (-1)^i R^{(i)} &= \sum_{i=0}^k (-1)^i \binom{k}{i} 3^i - \sum_{i=0}^k (-1)^i M^{(i)} \\ &= (-3+1)^k - \sum_{i=0}^k (-1)^i M^{(i)} \\ &= (-2)^k - \sum_{i=0}^k (-1)^i M^{(i)}, \end{aligned}$$

giving us the desired result. ■

At this point we can see how Conjecture 5.5.14 falls naturally out of our arithmetic.

Conjecture 6.5.9 (Walford's Conjecture¹) In the scenario described in Definitions 6.5.5, the alternating sum of partial missing region counts is bounded by

$$\sum_{i=0}^k (-1)^{k-i} M^{(i)} \geq 0,$$

with equality if and only if $M^{(0)} = M^{(1)} = \dots = M^{(k)} = 0$.

If this conjecture were true we could apply it to Corollary 6.5.8 to obtain

$$(-1)^k \sigma_k \leq 2^k + \begin{cases} (-1)^{d-k+1} & \text{if } 2k \geq d; \\ 0 & \text{if } 2k < d, \end{cases}$$

leading us through the use of Lemma 6.3.4 to the following result.

¹Named after Philip Walford, a fellow student at the American Institute of Mathematics who was going to prove this for me in *his* thesis.

Consequence 6.5.10 Consider a collection of k balanced $(d-1)$ -fans in a convex d -dimensional space \mathcal{P} as described above. Assume that k is as large as possible, i.e., $k = d-1$. If $\Phi > 1$ then the maximal face count Φ is bounded above by

$$\Phi \leq 2^{d-1} + 1,$$

the precise bound suggested by Conjecture 5.5.14.

Conjecture 6.5.9 is trivial to prove for very small d , as seen below.

Lemma 6.5.11 Conjecture 6.5.9 is true for $d \leq 4$.

Proof If $d \leq 4$ then the condition $2h \leq d$ is true for all $h \leq d-2$. Since $k \leq d-1$, it follows from Corollary 6.5.7 that $M^{(0)} = M^{(1)} = \dots = M^{(k-1)} = 0$ and so

$$\sum_{i=0}^k (-1)^{k-i} M^{(i)} = M^{(k)} \geq 0,$$

with equality if and only if $M^{(k)} = 0$. ■

As d increases however, the verification of Conjecture 6.5.9 requires further conditions upon the relationships between missing regions.

Definition 6.5.12 (Conjunction of Vectors) In the scenario described in Definitions 6.5.4, suppose that our k fans are partitioned into two disjoint sets S and T (so that $\{\mathcal{F}_1, \dots, \mathcal{F}_k\} = S \cup T$).

Then for any $\mathbf{u} \in \mathfrak{R}(S)$ and $\mathbf{v} \in \mathfrak{R}(T)$ we define the *conjunction* $\mathbf{u} \star \mathbf{v}$ to be the vector in the overall region space $\mathfrak{R}(S \cup T)$ whose associated regions are the associated regions of \mathbf{u} and the associated regions of \mathbf{v} .

For instance, if $k = 5$ and we consider vectors $\mathbf{u} = (i_2, i_4, i_5) \in \mathfrak{R}(\mathcal{F}_2, \mathcal{F}_4, \mathcal{F}_5)$ and $\mathbf{v} = (i_1, i_3) \in \mathfrak{R}(\mathcal{F}_1, \mathcal{F}_3)$ then their conjunction is $\mathbf{u} \star \mathbf{v} = (i_1, i_2, i_3, i_4, i_5) \in \mathfrak{R}(\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3, \mathcal{F}_4, \mathcal{F}_5)$ with associated regions $\rho_{1,i_1}, \dots, \rho_{5,i_5}$.

Conjunction offers a natural way of expressing $\mathfrak{R}(S \cup T)$ as the sum $\mathfrak{R}(S) \oplus \mathfrak{R}(T)$ by splitting out $|S|$ of the coordinate positions of \mathbb{Z}_3^k into $\mathfrak{R}(S)$ and the remaining $|T|$ coordinate positions into $\mathfrak{R}(T)$. In particular we see that each vector of $\mathfrak{R}(S \cup T)$ has a unique representation as $\mathbf{u} \star \mathbf{v}$ for $\mathbf{u} \in \mathfrak{R}(S)$ and $\mathbf{v} \in \mathfrak{R}(T)$.

Lemma 6.5.13 Consider the scenario described in Definitions 6.5.4, and let our k fans be partitioned into two disjoint sets S and T . Then for any $\mathbf{u} \in \mathfrak{R}(S)$, \mathbf{u} is a missing region in $\mathfrak{R}(S)$ if and only if $\mathbf{u} \star \mathbf{v}$ is a missing region in $\mathfrak{R}(S \cup T)$ for each $\mathbf{v} \in \mathfrak{R}(T)$.

Proof Without loss of generality, let $S = \{\mathcal{F}_1, \dots, \mathcal{F}_s\}$ and $T = \{\mathcal{F}_{s+1}, \dots, \mathcal{F}_k\}$. Let $\mathbf{u} = (u_1, \dots, u_s)$.

If \mathbf{u} is a missing region in $\mathfrak{R}(S)$ then the associated regions of \mathbf{u} have empty intersection, i.e., $\rho_{1,u_1} \cap \dots \cap \rho_{s,u_s} = \emptyset$. This implies that for any $\mathbf{v} = (v_{s+1}, \dots, v_k) \in \mathfrak{R}(T)$ we also have $\rho_{1,u_1} \cap \dots \cap \rho_{s,u_s} \cap \rho_{s+1,v_{s+1}} \cap \dots \cap \rho_{k,v_k} = \emptyset$, and so $\mathbf{u} \star \mathbf{v}$ is a missing region in $\mathfrak{R}(S \cup T)$.

Conversely, suppose that $\mathbf{u} \star \mathbf{v}$ is a missing region for each $\mathbf{v} \in \mathfrak{R}(T)$. Recall from Definitions 6.5.4 that $\rho_{i,0} \cup \rho_{i,1} \cup \rho_{i,2} = \mathcal{P} \setminus \mathcal{F}_i$ for each i . Examining the closures $\overline{\rho_{i,j}}$ it follows that $\overline{\rho_{i,0}} \cup \overline{\rho_{i,1}} \cup \overline{\rho_{i,2}} = \overline{\mathcal{P} \setminus \mathcal{F}_i} = \mathcal{P}$ for each i , since a union of closures is the closure of the union. Thus we find that

$$\begin{aligned} \overline{\rho_{1,u_1}} \cap \dots \cap \overline{\rho_{s,u_s}} &= \overline{\rho_{1,u_1}} \cap \dots \cap \overline{\rho_{s,u_s}} \cap \left(\bigcap_{i=s+1}^k (\overline{\rho_{i,0}} \cup \overline{\rho_{i,1}} \cup \overline{\rho_{i,2}}) \right) \\ &= \overline{\rho_{1,u_1}} \cap \dots \cap \overline{\rho_{s,u_s}} \cap \left(\bigcup_{(v_{s+1}, \dots, v_k) \in \mathfrak{R}(T)} (\overline{\rho_{s+1,v_{s+1}}} \cap \dots \cap \overline{\rho_{k,v_k}}) \right) \\ &= \bigcup_{(v_{s+1}, \dots, v_k) \in \mathfrak{R}(T)} (\overline{\rho_{1,u_1}} \cap \dots \cap \overline{\rho_{s,u_s}} \cap \overline{\rho_{s+1,v_{s+1}}} \cap \dots \cap \overline{\rho_{k,v_k}}). \end{aligned} \quad (6.17)$$

But since $\mathbf{u} \star \mathbf{v}$ is a missing region for each $\mathbf{v} \in \mathfrak{R}(T)$, we know for each $(v_{s+1}, \dots, v_k) \in \mathfrak{R}(T)$ that $\rho_{1,u_1} \cap \dots \cap \rho_{s,u_s} \cap \rho_{s+1,v_{s+1}} \cap \dots \cap \rho_{k,v_k} = \emptyset$ and so (since the interior of an intersection is the intersection of interiors) we have

$$\text{int}(\overline{\rho_{1,u_1}} \cap \dots \cap \overline{\rho_{s,u_s}} \cap \overline{\rho_{s+1,v_{s+1}}} \cap \dots \cap \overline{\rho_{k,v_k}}) = \emptyset.$$

Thus from Equation 6.17 we see that $\overline{\rho_{1,u_1}} \cap \dots \cap \overline{\rho_{s,u_s}}$ is a finite union of convex spaces with empty interior and so

$$\text{int}(\overline{\rho_{1,u_1}} \cap \dots \cap \overline{\rho_{s,u_s}}) = \emptyset.$$

Observing again that the interior of an intersection is the intersection of interiors, it follows that $\rho_{1,u_1} \cap \dots \cap \rho_{s,u_s} = \emptyset$ and so \mathbf{u} is a missing region in $\mathfrak{R}(S)$. ■

Lemma 6.5.14 Again consider the scenario described in Definitions 6.5.4, and let our k fans be partitioned into two disjoint sets S and T for which $S \neq \emptyset$ and $2|T| < d$. Then there are two vectors $\mathbf{u}, \mathbf{u}' \in \mathfrak{R}(S)$ for which the following properties hold.

- Vectors \mathbf{u} and \mathbf{u}' have all coordinates distinct. That is, if $\mathbf{u} = (x_1, \dots, x_s)$ and $\mathbf{u}' = (y_1, \dots, y_s)$ then $x_i \neq y_i$ for all i .
- Vectors $\mathbf{u} \star \mathbf{v}$ and $\mathbf{u}' \star \mathbf{v}$ are visible regions for all $\mathbf{v} \in \mathfrak{R}(T)$.

Proof Let the fans in T be $\mathcal{F}_1, \dots, \mathcal{F}_r$ with axes $\mathcal{A}_1, \dots, \mathcal{A}_r$ respectively, where $r = |T|$. By general position these r axes intersect in a $(d - 2r)$ -dimensional hyperplane \mathcal{H} containing central point \mathbf{o} in its relative interior, noting that $d - 2r > 0$. We can thus use a vector system in which \mathbf{o} is the origin.

Consider then any fan \mathcal{F} in S . From Lemma 6.2.4 the intersection of any non-strict leaf of \mathcal{F} with \mathcal{H} is either an original $(r, 1)$ -piece, an original $(r + 1, 0)$ -piece or the single point \mathbf{o} . Lemma 6.2.3 shows that in each of these situations the non-strict leaf intersects \mathcal{H} in at most a $(d - 2r - 1)$ -dimensional object. Since there are finitely many fans in S and a fan is the union of its three non-strict leaves, it follows that there is some point $\mathbf{p} \in \mathcal{H}$ for which both \mathbf{p} and $-\mathbf{p}$ belong to none of the fans in S .

Let the fans in S be $\mathcal{F}_{r+1}, \dots, \mathcal{F}_{r+s}$. It follows from the above result that for each j in the range $r+1 \leq j \leq r+s$, points \mathbf{p} and $-\mathbf{p}$ belong to open regions ρ_{j,x_j} and ρ_{j,y_j} respectively for some $x_j, y_j \in \mathbb{Z}_3$. Furthermore we must have $x_j \neq y_j$ for each j since no single open region ρ_{j,x_j} can contain both \mathbf{p} and $-\mathbf{p}$ (otherwise convexity would require this open region to contain \mathbf{o} also).

So defining $\mathbf{u} = (x_{r+1}, \dots, x_{r+s})$ and $\mathbf{u}' = (y_{r+1}, \dots, y_{r+s})$ we see that both \mathbf{u} and \mathbf{u}' are members of $\mathfrak{R}(S)$ for which the following properties hold.

- Vectors \mathbf{u} and \mathbf{u}' have all coordinates distinct, i.e., $x_j \neq y_j$ for each j ;
- A small neighbourhood of \mathbf{p} belongs to each associated region of \mathbf{u} and a small neighbourhood of $-\mathbf{p}$ belongs to each associated region of \mathbf{u}' .

We claim that these vectors $\mathbf{u}, \mathbf{u}' \in \mathfrak{R}(S)$ are the two vectors that we seek. All that remains to be shown is that $\mathbf{u} \star \mathbf{v}$ and $\mathbf{u}' \star \mathbf{v}$ are visible regions for all $\mathbf{v} \in \mathfrak{R}(T)$.

Choose any $\mathbf{v} = (z_1, \dots, z_r) \in \mathfrak{R}(T)$ and for each i let \mathcal{L}_i denote the strict leaf of fan \mathcal{F}_i that does not border region ρ_{i,z_i} . Since $2r < d$, we see from the proof of point 4 of Lemma 6.3.1 that every selection of $r-1$ axes and one leaf from r distinct fans gives rise to an original $(r-1, 1)$ -piece. In particular for each i in the range $1 \leq i \leq r$ there is some point \mathbf{q}_i belonging to strict leaf \mathcal{L}_i and each axis \mathcal{A}_j for $1 \leq j \leq r, j \neq i$.

Consider then the point $\mathbf{r} = \mathbf{p} - \varepsilon(\mathbf{q}_1 + \dots + \mathbf{q}_r)$, where $\varepsilon > 0$ is small enough to keep \mathbf{r} within the small neighbourhood of \mathbf{p} described above. For each i in the range $1 \leq i \leq r$ we see that $\varepsilon \mathbf{q}_i$ lies on strict leaf \mathcal{L}_i and so since our fans are balanced the point $-\varepsilon \mathbf{q}_i$ lies within the opposite region ρ_{i,z_i} . But for each $j \neq i$ point \mathbf{q}_j lies on the axis \mathcal{A}_i , and likewise $\mathbf{p} \in \mathcal{A}_i$. Thus \mathbf{r} lies in the region ρ_{i,z_i} .

Therefore \mathbf{r} lies in each associated region of \mathbf{v} , and from the above properties of \mathbf{u} we know that \mathbf{r} lies in each associated region of \mathbf{u} . Thus $\mathbf{u} \star \mathbf{v}$ is a visible region, and by a similar argument $\mathbf{u}' \star \mathbf{v}$ is similarly a visible region. ■

As a simple example of how Lemma 6.5.14 might be used, we can prove the following result.

Lemma 6.5.15 In the scenario described in Definitions 6.5.4, if $2(k-1) < d$ then there is at most one missing region, i.e., $M \leq 1$.

Proof Let our k fans be $\mathcal{F}_1, \dots, \mathcal{F}_k$ and set $S = \{\mathcal{F}_1\}$ and $T = \{\mathcal{F}_2, \mathcal{F}_3, \dots, \mathcal{F}_k\}$. Since $|T| = k-1$ we find that $2|T| < d$ and so we can apply Lemma 6.5.14.

In this case Lemma 6.5.14 implies that there are some distinct $x_1, y_1 \in \mathbb{Z}_3$ or equivalently $(x_1), (y_1) \in \mathfrak{R}(S) = \mathbb{Z}_3^1$ for which vectors $(x_1) \star \mathbf{v}$ and $(y_1) \star \mathbf{v}$ are visible regions for all $\mathbf{v} \in \mathfrak{R}(T)$. Defining z_1 to be the member of \mathbb{Z}_3 that is neither x_1 nor y_1 , it follows that every missing region in $\mathfrak{R}(S \cup T)$ must have first coordinate z_1 .

Similarly, for any j we can let $S = \{\mathcal{F}_j\}$ and use Lemma 6.5.14 to find some $z_j \in \mathbb{Z}_3$ for which every missing region in the overall region space \mathfrak{R} must have j th coordinate z_j . Combining these results for all $j = 1, 2, \dots, k$ we see that the only possible missing region in the overall region space is (z_1, z_2, \dots, z_k) . Thus $M \leq 1$. ■

Lemma 6.5.16 Conjecture 6.5.9 is true for $d \leq 5$.

Proof If $d \leq 4$ then Conjecture 6.5.9 is already proven in Lemma 6.5.11. So let $d = 5$. In this case the condition $2h \leq d$ is true for all $h \leq 2$.

Suppose that $k \leq 3$. Then $2h \leq d$ is true for all $h \leq k - 1$, and so Corollary 6.5.7 implies that $M^{(0)} = M^{(1)} = \dots = M^{(k-1)} = 0$ and again

$$\sum_{i=0}^k (-1)^{k-i} M^{(i)} = M^{(k)} \geq 0$$

with equality if and only if $M^{(k)} = 0$.

This leaves the case $k = d - 1 = 4$. Here we have $2h \leq d$ true for all $h \leq k - 2$ giving $M^{(0)} = M^{(1)} = \dots = M^{(k-2)} = 0$ and so

$$\sum_{i=0}^k (-1)^{k-i} M^{(i)} = M^{(k)} - M^{(k-1)}. \quad (6.18)$$

Thus we must examine $M^{(k-1)}$ in more detail. Let S be the set of all k fans; then $M^{(k-1)}$ counts the missing regions in $\mathfrak{R}(T)$ for all subsets $T \subseteq S$ of size $|T| = k - 1$.

For $(d, k) = (5, 4)$ we have $2(k - 2) < d$, and so for any such subset T of $k - 1$ fans we see from Lemma 6.5.15 that there is at most one missing region in $\mathfrak{R}(T)$. Since there are $\binom{k}{k-1} = k$ such subsets T it follows that $M^{(k-1)} \leq k = 4$.

From Lemma 6.5.13 we see that for any subset T of $k - 1$ fans, each missing region in $\mathfrak{R}(T)$ corresponds to a line of three missing regions in \mathfrak{R} . For instance, if (r_1, \dots, r_{k-1}) were a missing region in $\mathfrak{R}(\mathcal{F}_1, \dots, \mathcal{F}_{k-1})$ then (r_1, \dots, r_{k-1}, x) would be a missing region in $\mathfrak{R}(\mathcal{F}_1, \dots, \mathcal{F}_k)$ for each $x \in \mathbb{Z}_3$.

In particular, if $M^{(k-1)} = 1$ then this single missing region in $\mathfrak{R}(T)$ for some subset T of $k - 1$ fans corresponds to three missing regions in $\mathfrak{R}(S)$ and so $M^{(k)} \geq 3$. Thus the alternating sum in Equation 6.18 is strictly positive.

If $M^{(k-1)} \geq 2$ then we have at least two different subsets T_1 and T_2 of $k - 1$ fans with missing regions in both $\mathfrak{R}(T_1)$ and $\mathfrak{R}(T_2)$. Since $T_1 \neq T_2$, the corresponding two lines of missing regions in $\mathfrak{R}(S)$ are not parallel and so have at most one point in common. Thus there are at least $2 \cdot 3 - 1 = 5$ missing regions in $\mathfrak{R}(S)$ giving $M^{(k)} \geq 5$. But since we already have $M^{(k-1)} \leq 4$, we see once again that the sum in Equation 6.18 is strictly positive.

Finally, if $M^{(k-1)} = 0$ then the sum in Equation 6.18 reduces to $M^{(k)}$ which is ≥ 0 with equality if and only if $M^{(k)} = 0$ also. ■

We can continue to prove Conjecture 6.5.9 for increasing values of d through more and more convoluted combinatorial arguments and case analyses, but a general proof for all d remains elusive. As discussed in the opening comments of Chapter 5, this is the point at which the attempted proof of the upper bound $\Phi \leq 2^{d-1} + 1$ breaks down.

6.6 Future Directions

In conclusion, we discuss the ways in which work is still underway to prove Conjecture 5.5.14 which implies that any intersection of balanced $(d - 1)$ -fans in a convex d -dimensional polytope,

all of whose axes pass through some common interior point \mathbf{o} , has at most $2^{d-1} + 1$ maximal faces. Recall that once proven, this conjecture implies Consequence 5.5.15 which establishes the upper bound of $2^{n-1} + 1$ maximal embedded faces in the projective solution space.

In addition to the approaches discussed below, there is a vast body of work already published on the combinatorial structure of polytopes such as can be seen in [2], [6] and [48]. Although this theory of polytopes contains many powerful results, sadly none were found that directly assist with the specific problems at hand.

6.6.1 Proving Walford's Conjecture

It is still hoped that a general proof can be found for Conjecture 6.5.9, which is the remaining hurdle for the general position case with $k = d - 1$. Certainly there are more constraints upon the various missing region counts that can be derived from the geometry of k fans in a d -dimensional space than are actually being used.

A more complex variant of Lemma 6.5.14 can be established for a partition of the k fans into disjoint sets S and T for which $2|T| = d$. Other relationships can be established involving the situations in which a fan is removed or where the d -dimensional space is restricted to the $(d - 2)$ -dimensional axis of one of our fans, leading to the possibility of an inductive proof.

If these relationships prove insufficient to establish Conjecture 6.5.9, it might be necessary to extract further conditions from the fact that these intersecting fans were produced from a 3-manifold triangulation. That is, we might require tighter constraints upon the geometry of the projective edge weight solution space than those proven in Chapter 5. It is hoped however that this will not be necessary.

6.6.2 General Position with $k \leq d - 2$

Once Conjecture 6.5.9 is established, the task remains to bound the maximal face count Φ for $k \leq d - 2$. From Consequence 6.5.10 and Lemma 5.5.10 we see that Φ ranges from $2^k + 1$ for $k = d - 1$ down to 3^k for $k \leq \frac{d}{2}$, raising the possibility that we might not easily find an exact bound for k within this range.

Although Corollary 6.5.8 and Conjecture 6.5.9 still produce the bound $|\sigma_k| \leq 2^k + 1$, we no longer have the convenience of knowing that $|\sigma_k| = \Phi$. Instead, whereas σ_k is an alternating sum of piece counts $G_{r,k-r}$, for $k \leq d - 2$ we find from Corollary 6.2.9 that Φ is bounded above by an absolute sum of piece counts $G_{r,k-r}$.

As described in the remarks following Corollary 6.2.9, the bound $\Phi \leq \sum_r G_{r,k-r}$ is quite loose. Nevertheless this bound is precise for $k = d - 1$ (since $\Phi = G_{0,k} = \sum_r G_{r,k-r}$). At the other end of the spectrum, for $2k < d$ we can calculate this sum precisely using Lemma 6.3.1 as

$$\sum_r G_{r,k-r} = \sum_{r=0}^k \binom{k}{r} 3^{k-r} = 4^k.$$

We see that $4^k = 2^{2k} \leq 2^{d-1}$, which remains below the bound suggested in Conjecture 5.5.14.

For the one remaining case in which Lemma 6.3.1 allows us to calculate precise values, i.e., $2k = d$, we find that $\sum_r G_{r,k-r} = 4^k - 1 = 2^d - 1$ which is approximately twice the bound

suggested in Conjecture 5.5.14. Nevertheless there is hope that we can refine Corollary 6.2.9 to produce a tighter bound for Φ in terms of the various piece counts $G_{r,k-r}$ that might lead us to the desired bound $\Phi \leq 2^{d-1} + 1$.

The problem still remains however of calculating the sum (or whatever other refined combination we might produce) of the piece counts $G_{r,k-r}$. With this aim in mind, we can generalise the Euler characteristic equations seen in Lemma 6.3.5 to a much larger family of equations of which Lemma 6.3.5 represents a special case. Specifically, for any non-negative r_0 and s_0 for which $r_0 + s_0 \leq k$ and $2r_0 + s_0 < d$ it can be shown that

$$(-1)^{s_0} G_{r_0, s_0} = \sum_{\substack{r+s \geq r_0+s_0 \\ 2r+s < d \\ r \geq r_0}} (-1)^s \binom{r}{r_0} \binom{s}{s_0} F_{r,s} + \begin{cases} (-1)^d \binom{k}{d} & \text{if } s_0 = 0 \text{ and } d \leq 2k; \\ 0 & \text{otherwise.} \end{cases} \quad (6.19)$$

Just as the proof of Lemma 6.3.5 involves examining the Euler characteristic of the intersection of h fans at a time, the proof of these more general relations involves calculating the Euler characteristic of the intersection of h fans at a time restricted to portions involving $\geq r_0$ axes and then manipulating the resulting equations to extract a single G_{r_0, s_0} term. Lemma 6.3.5 can then be reestablished by summing Equation 6.19 for all r_0 and s_0 for which $r_0 + s_0 = h$.

It is thus possible that, just as we have an expansion of σ_h in terms of the τ_i (Lemma 6.3.5) which is then converted into an expansion of τ_h in terms of the σ_i (Lemma 6.3.8), we might likewise take the above expansion of G_{r_0, s_0} in terms of the $F_{r,s}$ and convert it into an expansion of F_{r_0, s_0} in terms of the $G_{r,s}$. We can then continue to mirror the algebraic manipulations of Section 6.4 until with luck we have an expression of Φ in terms of a small number of h -partial final piece counts that are easy to deal with, such as the h -partial region counts $F_{0,0}^{(h)}$.

There is a final approach to the case $k \leq d-2$ which appears to offer some promise, and this is to first deal with the case $k = d-1$ in non-general position as discussed in the following section. We then make the simple observation that any intersection of k fans for $k \leq d-2$ is in fact an intersection of $d-1$ fans in which some of the fans are duplicated, and so for $k \leq d-2$ there is no more work to be done.

6.6.3 Non-General Position

All the work that is presented in Chapter 6 relies upon the rather grand assumption that our k fans in our d -dimensional polytope are in general position. In examining real examples of 3-manifold triangulations and their edge weight solution spaces, one finds that this is frequently not the case. However, as discussed in the closing remarks of Chapter 5, it is still felt that for $k \leq d-1$ the number of maximal faces is largest in the general position scenario. The strategy for dealing with the more general unrestricted problem is then as follows.

Using Definition 6.1.1 we can construct a measure of how far removed from general position any collection of fans is. For example, we might measure how many different intersections of leaves and/or axes break the specific requirements that Definition 6.1.1 sets upon them. Given any set of $k \leq d-1$ fans that are not in general position, we then provide a method of adjusting one of these fans slightly so that the resulting collection of fans is “closer” to general position according to the measure previously defined. Through discreteness arguments we show that this procedure

can only be repeated finitely many times before we have a collection of k fans that are indeed in general position.

The key to the argument then will be the way in which this fan adjustment takes place. Specifically, the adjustment must be done in such a way that the number of maximal faces of the overall intersection is not reduced. In this way we prove that the largest number of maximal faces is indeed seen in the general position scenario, from which point the material of Chapter 6 can be used to bound the maximal face count.

But is it always possible to make such an adjustment? Intuitively what characterises a non-general-position scenario is that certain leaves and axes from different fans that should intersect to form a t -dimensional object in fact either intersect to form a larger dimensional object or do not intersect at all. It feels then that when moving closer to general position, it should be possible to simply reduce a larger dimensional intersection to one or more smaller dimensional intersections, and in this way reduce each maximal face to one or more smaller maximal faces so that our overall maximal face count is not diminished. Thus such a method of proof does appear to have potential.

However, this method of adjustment breaks down when our maximal faces are 1-dimensional; in this case reducing them to 0-dimensional intersections simply reduces them to the central point **o**. Recalling from Lemma 5.5.2 that each maximal face can be expressed as an intersection of leaves and that each leaf is a $(d - 1)$ -dimensional half-hyperplane, we can specifically expect this method to break down when $k = d$. This expectation is consistent with Lemma 6.1.4 and the closing remarks of Chapter 5.

It is known from Theorem 5.5.1 however that we might indeed have $k = d$ and so this case will need to be handled separately. Experimental results for small d suggest that all intersections of d fans in a d -dimensional space remain true to the bound $\Phi \leq 2^{d-1} + 1$. Work towards a proof of this claim is still continuing.

Bibliography

- [1] Colin C. Adams, *The knot book: An elementary introduction to the mathematical theory of knots*, W. H. Freeman & Co., New York, 1994.
- [2] Arne Brøndsted, *An introduction to convex polytopes*, Graduate Texts in Mathematics, no. 90, Springer-Verlag, New York, 1983.
- [3] Benjamin A. Burton, *Regina (Normal surface and 3-manifold software)*, <http://regina.sourceforge.net/>, 1999–2003.
- [4] Patrick J. Callahan, John C. Dean, and Jeffrey R. Weeks, *The simplest hyperbolic knots*, J. Knot Theory Ramifications **8** (1999), no. 3, 279–297.
- [5] Patrick J. Callahan, Martin V. Hildebrand, and Jeffrey R. Weeks, *A census of cusped hyperbolic 3-manifolds*, Math. Comp. **68** (1999), no. 225, 321–332.
- [6] Günter Ewald, *Combinatorial convexity and algebraic geometry*, Graduate Texts in Mathematics, no. 168, Springer, New York, 1996.
- [7] Mark Giesbrecht, *Fast computation of the Smith normal form of an integer matrix*, Proceedings of the 1995 International Symposium on Symbolic and Algebraic Computation (Montreal, Quebec, Canada), ACM Press, 1995, pp. 110–118.
- [8] Wolfgang Haken, *Theorie der Normalflächen*, Acta Math. **105** (1961), 245–375.
- [9] ———, *Über das Homöomorphieproblem der 3-Mannigfaltigkeiten. I*, Math. Z. **80** (1962), 89–120.
- [10] Geoffrey Hemion, *The classification of knots and 3-dimensional spaces*, Oxford Science Publications, Oxford University Press, Oxford, 1992.
- [11] John Hempel, *3-manifolds*, Annals of Mathematics Studies, no. 86, Princeton University Press, Princeton, NJ, 1976.
- [12] Martin V. Hildebrand and Jeffrey R. Weeks, *A computer generated census of cusped hyperbolic 3-manifolds*, Computers and Mathematics (Cambridge, MA, 1989), Springer, New York, 1989, pp. 53–59.
- [13] Heinz Hopf, *Zum Clifford-Kleinschen Raumproblem*, Math. Ann. **95** (1925–26), 313–319.

- [14] William Jaco, *Lectures on three-manifold topology*, Regional Conference Series in Mathematics, no. 43, Amer. Math. Soc., Providence, RI, 1980.
- [15] William Jaco, David Letscher, and J. Hyam Rubinstein, *Algorithms for essential surfaces in 3-manifolds*, Topology and Geometry: Commemorating SISTAG, Contemporary Mathematics, no. 314, Amer. Math. Soc., Providence, RI, 2002, pp. 107–124.
- [16] William Jaco and Ulrich Oertel, *An algorithm to decide if a 3-manifold is a Haken manifold*, Topology **23** (1984), no. 2, 195–209.
- [17] William Jaco and J. Hyam Rubinstein, *0-efficient triangulations of 3-manifolds*, Preprint, July 2002.
- [18] ———, *1-efficient triangulations of 3-manifolds*, In preparation, 2003.
- [19] William Jaco and Jeffrey L. Tollefson, *Algorithms for the complete decomposition of a closed 3-manifold*, Illinois J. Math. **39** (1995), no. 3, 358–406.
- [20] Paik Kee Kim, *Some 3-manifolds which admit Klein bottles*, Trans. Amer. Math. Soc. **244** (1978), 299–312.
- [21] Hellmuth Kneser, *Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten*, Jahresbericht der Deut. Math. Verein. **38** (1929), 248–260.
- [22] Marc Lackenby, *Word hyperbolic Dehn surgery*, Invent. Math. **140** (2000), no. 2, 243–282.
- [23] ———, *Taut ideal triangulations of 3-manifolds*, Geom. Topol. **4** (2000), 369–395 (electronic).
- [24] Bruno Martelli and Carlo Petronio, *Three-manifolds having complexity at most 9*, Experiment. Math. **10** (2001), no. 2, 207–236.
- [25] ———, *A new decomposition theorem for 3-manifolds*, Illinois J. Math. **46** (2002), 755–780.
- [26] ———, *Complexity of geometric three-manifolds*, Preprint, March 2003.
- [27] William S. Massey, *A basic course in algebraic topology*, Graduate Texts in Mathematics, no. 127, Springer-Verlag, New York, 1991.
- [28] Sergei V. Matveev, *Computer recognition of three-manifolds*, Experiment. Math. **7** (1998), no. 2, 153–161.
- [29] ———, *Tables of 3-manifolds up to complexity 6*, Max-Planck-Institut für Mathematik Preprint Series (1998), no. 67, available from <http://www.mpim-bonn.mpg.de/html/preprints/preprints.html>.
- [30] ———, *Computer classification of 3-manifolds*, Russ. J. Math. Phys. **7** (2000), no. 3, 319–329.
- [31] John Milnor, *Groups which act on S^n without fixed points*, Amer. J. Math. **79** (1957), 623–630.
- [32] Peter Orlik, *Seifert manifolds*, Lecture Notes in Mathematics, no. 291, Springer-Verlag, Berlin, 1972.

- [33] Peter Orlik and Hiroaki Terao, *Arrangements of hyperplanes*, Grundlehren der mathematischen Wissenschaften, no. 300, Springer-Verlag, Berlin, 1992.
- [34] Udo Pachner, *P.L. homeomorphic manifolds are equivalent by elementary shellings*, European J. Combin. **12** (1991), no. 2, 129–145.
- [35] Georg Pick, *Geometrisches zur Zahlenlehre*, Naturwiss. Z. Lotos (Prague) **19** (1899), 311–319.
- [36] Igor Rivin, *Euclidean structures on simplicial surfaces and hyperbolic volume*, Ann. of Math. (2) **139** (1994), no. 3, 553–580.
- [37] ———, *Combinatorial optimization in geometry*, Preprint, July 1999.
- [38] J. Hyam Rubinstein, *On 3-manifolds that have finite fundamental group and contain Klein bottles*, Trans. Amer. Math. Soc. **251** (1979), 129–137.
- [39] ———, *An algorithm to recognize the 3-sphere*, Proceedings of the International Congress of Mathematicians (Zürich, 1994), vol. 1, Birkhäuser, 1995, pp. 601–611.
- [40] ———, *Polyhedral minimal surfaces, Heegaard splittings and decision problems for 3-dimensional manifolds*, Geometric Topology (Athens, GA, 1993), AMS/IP Stud. Adv. Math., vol. 2, Amer. Math. Soc., 1997, pp. 1–20.
- [41] J. Hyam Rubinstein and Joan S. Birman, *One-sided Heegaard splittings and homeotopy groups of some 3-manifolds*, Proc. London Math. Soc. (3) **49** (1984), no. 3, 517–536.
- [42] Horst Schubert, *Bestimmung der Primfaktorzerlegung von Verkettungen*, Math. Z. **76** (1961), 116–148.
- [43] Arne Storjohann, *Near optimal algorithms for computing Smith normal forms of integer matrices*, Proceedings of the 1996 International Symposium on Symbolic and Algebraic Computation (Zürich, Switzerland), ACM Press, 1996, pp. 267–274.
- [44] Stephan Tillmann, *On character varieties : surfaces associated to mutation & deformation of hyperbolic 3-manifolds*, Ph.D. thesis, University of Melbourne, 2002.
- [45] Jeffrey L. Tollefson, *Normal surface Q-theory*, Pacific J. Math. **183** (1998), no. 2, 359–374.
- [46] Vladimir G. Turaev and Oleg Y. Viro, *State sum invariants of 3-manifolds and quantum 6j-symbols*, Topology **31** (1992), no. 4, 865–902.
- [47] Jeffrey R. Weeks, *Snappea (Hyperbolic 3-manifold software)*, <http://www.northnet.org/weeks/index/SnapPea.html>, 1991–2000.
- [48] Günter M. Ziegler, *Lectures on polytopes*, Graduate Texts in Mathematics, no. 152, Springer-Verlag, New York, 1995.

Appendix A

Tables of Triangulations

Enumerated in the following tables are all closed prime minimal triangulations formed from at most six tetrahedra, both orientable and non-orientable. These triangulations are arranged according to their underlying 3-manifolds. For each 3-manifold we present the following information.

- Δ : The number tetrahedra in each minimal triangulation;
- *3-Manifold*: The common name of the 3-manifold as described in Section 3.1;
- *Mat.*: The index assigned to the 3-manifold in the orientable census of Matveev [29];
- *Seifert Structure*: The Seifert fibred space parameters as described in Section 3.1.2;
- *Triangulations*: A list of all minimal triangulations of the 3-manifold as described in Sections 3.4 and 3.5;
- *Homology*: The first homology group of the 3-manifold.

Note that the Matveev index and the Seifert structure are presented only for orientable 3-manifolds. Furthermore the Seifert structure for Lens spaces is well known and thus omitted. Each 3-manifold of the form $K^2 \tilde{\times} I \cup K^2 \tilde{\times} I / \dots$ is written using the abbreviated notation $(K^2 \tilde{\times} I)^2 / \dots$ for space considerations.

Table A.1: Closed orientable prime minimal triangulations

Δ	3-Manifold	Mat.	Seifert Structure	Triangulations	Homology
1	S^3	0_1		$L_{1,0}, C_1$	0
	$L(4, 1)$	1_1		$L_{4,1}$	\mathbb{Z}_4
	$L(5, 2)$	1_2		$L_{5,2}$	\mathbb{Z}_5
2	$S^2 \times S^1$	–		$L_{0,1}$	\mathbb{Z}
	$\mathbb{R}P^3$	0_2		$L_{2,1}, C_2$	\mathbb{Z}_2
	$L(3, 1)$	0_3		$L_{3,1}^{(1)}, L_{3,1}^{(2)}, L'_{3,1}$	\mathbb{Z}_3
	$L(5, 1)$	2_1		$L_{5,1}$	\mathbb{Z}_5

Δ	3-Manifold	Mat.	Seifert Structure	Triangulations	Homology
2	$L(7, 2)$	2_2		$L_{7,2}$	\mathbb{Z}_7
	$L(8, 3)$	2_3		$L_{8,3}$	\mathbb{Z}_8
	S^3/Q_8	2_4		\tilde{C}_2	$\mathbb{Z}_2 \oplus \mathbb{Z}_2$
3	$L(6, 1)$	3_1		$L_{6,1}$	\mathbb{Z}_6
	$L(9, 2)$	3_2		$L_{9,2}$	\mathbb{Z}_9
	$L(10, 3)$	3_3		$L_{10,3}$	\mathbb{Z}_{10}
	$L(11, 3)$	3_4		$L_{11,3}$	\mathbb{Z}_{11}
	$L(12, 5)$	3_5		$L_{12,5}$	\mathbb{Z}_{12}
	$L(13, 5)$	3_6		$L_{13,5}$	\mathbb{Z}_{13}
	S^3/Q_{12}	3_7	$S^2 : (2, 1) (2, 1) (3, -2)$	\tilde{C}_3	\mathbb{Z}_4
4	$L(7, 1)$	4_1		$L_{7,1}$	\mathbb{Z}_7
	$L(11, 2)$	4_2		$L_{11,2}$	\mathbb{Z}_{11}
	$L(13, 3)$	4_3		$L_{13,3}$	\mathbb{Z}_{13}
	$L(14, 3)$	4_4		$L_{14,3}$	\mathbb{Z}_{14}
	$L(15, 4)$	4_5		$L_{15,4}$	\mathbb{Z}_{15}
	$L(16, 7)$	4_6		$L_{16,7}$	\mathbb{Z}_{16}
	$L(17, 5)$	4_7		$L_{17,5}$	\mathbb{Z}_{17}
	$L(18, 5)$	4_8		$L_{18,5}$	\mathbb{Z}_{18}
	$L(19, 7)$	4_9		$L_{19,7}$	\mathbb{Z}_{19}
	$L(21, 8)$	4_{10}		$L_{21,8}$	\mathbb{Z}_{21}
	$S^3/Q_8 \times \mathbb{Z}_3$	4_{11}	$S^2 : (2, 1) (2, 1) (2, 1)$	$A_{2,-3}$	$\mathbb{Z}_2 \oplus \mathbb{Z}_6$
	S^3/Q_{16}	4_{12}	$S^2 : (2, 1) (2, 1) (4, -3)$	\tilde{C}_4	$\mathbb{Z}_2 \oplus \mathbb{Z}_2$
	S^3/D_{24}	4_{13}	$S^2 : (2, 1) (2, 1) (3, -1)$	$A_{3,-2}, J_{1 2,-1}$	\mathbb{Z}_8
	S^3/P_{24}	4_{14}	$S^2 : (2, 1) (3, 1) (3, -2)$	$C_{2,2}$	\mathbb{Z}_3
5	$L(8, 1)$	5_1		$L_{8,1}$	\mathbb{Z}_8
	$L(13, 2)$	5_2		$L_{13,2}$	\mathbb{Z}_{13}
	$L(16, 3)$	5_3		$L_{16,3}$	\mathbb{Z}_{16}
	$L(17, 3)$	5_4		$L_{17,3}$	\mathbb{Z}_{17}
	$L(17, 4)$	5_5		$L_{17,4}$	\mathbb{Z}_{17}
	$L(19, 4)$	5_6		$L_{19,4}$	\mathbb{Z}_{19}
	$L(20, 9)$	5_7		$L_{20,9}$	\mathbb{Z}_{20}
	$L(22, 5)$	5_8		$L_{22,5}$	\mathbb{Z}_{22}

Δ	3-Manifold	Mat.	Seifert Structure	Triangulations	Homology
5	$L(23, 5)$	5 ₉		$L_{23,5}$	\mathbb{Z}_{23}
	$L(23, 7)$	5 ₁₀		$L_{23,7}$	\mathbb{Z}_{23}
	$L(24, 7)$	5 ₁₁		$L_{24,7}$	\mathbb{Z}_{24}
	$L(25, 7)$	5 ₁₂		$L_{25,7}$	\mathbb{Z}_{25}
	$L(25, 9)$	5 ₁₃		$L_{25,9}$	\mathbb{Z}_{25}
	$L(26, 7)$	5 ₁₄		$L_{26,7}$	\mathbb{Z}_{26}
	$L(27, 8)$	5 ₁₅		$L_{27,8}$	\mathbb{Z}_{27}
	$L(29, 8)$	5 ₁₆		$L_{29,8}$	\mathbb{Z}_{29}
	$L(29, 12)$	5 ₁₇		$L_{29,12}$	\mathbb{Z}_{29}
	$L(30, 11)$	5 ₁₈		$L_{30,11}$	\mathbb{Z}_{30}
	$L(31, 12)$	5 ₁₉		$L_{31,12}$	\mathbb{Z}_{31}
	$L(34, 13)$	5 ₂₀		$L_{34,13}$	\mathbb{Z}_{34}
	$S^3/Q_8 \times \mathbb{Z}_5$	5 ₂₁	$S^2 : (2, 1) (2, 1) (2, 3)$	$A_{2,-5}, A_{2,-3 2,-3}$	$\mathbb{Z}_2 \oplus \mathbb{Z}_{10}$
	$S^3/Q_{12} \times \mathbb{Z}_5$	5 ₂₂	$S^2 : (2, 1) (2, 1) (3, 2)$	$A_{3,-5}, A_{2,-3 3,-2}$	\mathbb{Z}_{20}
	$S^3/Q_{16} \times \mathbb{Z}_3$	5 ₂₃	$S^2 : (2, 1) (2, 1) (4, -1)$	$A_{4,-3}, J_{1 3,-1}$	$\mathbb{Z}_2 \oplus \mathbb{Z}_6$
	S^3/Q_{20}	5 ₂₄	$S^2 : (2, 1) (2, 1) (5, -4)$	\tilde{C}_5	\mathbb{Z}_4
	$S^3/Q_{20} \times \mathbb{Z}_3$	5 ₂₅	$S^2 : (2, 1) (2, 1) (5, -2)$	$A_{5,-3}, J_{1 3,-2}$	\mathbb{Z}_{12}
	S^3/D_{40}	5 ₂₆	$S^2 : (2, 1) (2, 1) (5, -3)$	$A_{5,-2}, X_{2 2,-1}, J_{1 2,-3}$	\mathbb{Z}_8
	S^3/D_{48}	5 ₂₇	$S^2 : (2, 1) (2, 1) (3, 1)$	$A_{3,-4}, A_{2,-3 3,-1}$	\mathbb{Z}_{16}
	$S^3/P_{24} \times \mathbb{Z}_5$	5 ₂₈	$S^2 : (2, 1) (3, 2) (3, -1)$	$A_{3,-2 3,-2}$	\mathbb{Z}_{15}
	S^3/P_{48}	5 ₂₉	$S^2 : (2, 1) (3, 1) (4, -3)$	$C_{2,3}, P'_0$	\mathbb{Z}_2
	S^3/P'_{72}	5 ₃₀	$S^2 : (2, 1) (3, 1) (3, -1)$	$A_{3,-2 3,-1}, J_{2 2,-1}$	\mathbb{Z}_9
	S^3/P_{120}	5 ₃₁	$S^2 : (2, 1) (3, 1) (5, -4)$	P_0	0
6	$L(9, 1)$	6 ₁		$L_{9,1}$	\mathbb{Z}_9
	$L(15, 2)$	6 ₂		$L_{15,2}$	\mathbb{Z}_{15}
	$L(19, 3)$	6 ₃		$L_{19,3}$	\mathbb{Z}_{19}
	$L(20, 3)$	6 ₄		$L_{20,3}$	\mathbb{Z}_{20}
	$L(21, 4)$	6 ₅		$L_{21,4}$	\mathbb{Z}_{21}
	$L(23, 4)$	6 ₆		$L_{23,4}$	\mathbb{Z}_{23}
	$L(24, 5)$	6 ₇		$L_{24,5}$	\mathbb{Z}_{24}
	$L(24, 11)$	6 ₈		$L_{24,11}$	\mathbb{Z}_{24}
	$L(27, 5)$	6 ₉		$L_{27,5}$	\mathbb{Z}_{27}

Δ	3-Manifold	Mat.	Seifert Structure	Triangulations	Homology
6	$L(28, 5)$	6_{10}		$L_{28,5}$	\mathbb{Z}_{28}
	$L(29, 9)$	6_{11}		$L_{29,9}$	\mathbb{Z}_{29}
	$L(30, 7)$	6_{12}		$L_{30,7}$	\mathbb{Z}_{30}
	$L(31, 7)$	6_{13}		$L_{31,7}$	\mathbb{Z}_{31}
	$L(31, 11)$	6_{14}		$L_{31,11}$	\mathbb{Z}_{31}
	$L(32, 7)$	6_{15}		$L_{32,7}$	\mathbb{Z}_{32}
	$L(33, 7)$	6_{16}		$L_{33,7}$	\mathbb{Z}_{33}
	$L(33, 10)$	6_{17}		$L_{33,10}$	\mathbb{Z}_{33}
	$L(34, 9)$	6_{18}		$L_{34,9}$	\mathbb{Z}_{34}
	$L(35, 8)$	6_{19}		$L_{35,8}$	\mathbb{Z}_{35}
	$L(36, 11)$	6_{20}		$L_{36,11}$	\mathbb{Z}_{36}
	$L(37, 8)$	6_{21}		$L_{37,8}$	\mathbb{Z}_{37}
	$L(37, 10)$	6_{22}		$L_{37,10}$	\mathbb{Z}_{37}
	$L(39, 14)$	6_{23}		$L_{39,14}$	\mathbb{Z}_{39}
	$L(39, 16)$	6_{24}		$L_{39,16}$	\mathbb{Z}_{39}
	$L(40, 11)$	6_{25}		$L_{40,11}$	\mathbb{Z}_{40}
	$L(41, 11)$	6_{26}		$L_{41,11}$	\mathbb{Z}_{41}
	$L(41, 12)$	6_{27}		$L_{41,12}$	\mathbb{Z}_{41}
	$L(41, 16)$	6_{28}		$L_{41,16}$	\mathbb{Z}_{41}
	$L(43, 12)$	6_{29}		$L_{43,12}$	\mathbb{Z}_{43}
	$L(44, 13)$	6_{30}		$L_{44,13}$	\mathbb{Z}_{44}
	$L(45, 19)$	6_{31}		$L_{45,19}$	\mathbb{Z}_{45}
	$L(46, 17)$	6_{32}		$L_{46,17}$	\mathbb{Z}_{46}
	$L(47, 13)$	6_{33}		$L_{47,13}$	\mathbb{Z}_{47}
	$L(49, 18)$	6_{34}		$L_{49,18}$	\mathbb{Z}_{49}
	$L(50, 19)$	6_{35}		$L_{50,19}$	\mathbb{Z}_{50}
	$L(55, 21)$	6_{36}		$L_{55,21}$	\mathbb{Z}_{55}
	$S^3/Q_8 \times \mathbb{Z}_7$	6_{37}	$S^2 : (2, 1) (2, 1) (2, 5)$	$A_{2,-7}, A_{2,-5 2,-3},$ $A_{2,-3 2,-3 2,-3}$	$\mathbb{Z}_2 \oplus \mathbb{Z}_{14}$
	$S^3/Q_{12} \times \mathbb{Z}_7$	6_{38}	$S^2 : (2, 1) (2, 1) (3, 4)$	$A_{3,-7}, A_{2,-5 3,-1},$ $A_{2,-3 3,-4}, A_{2,-3 2,-3 3,-1}$	\mathbb{Z}_{28}
	$S^3/Q_{16} \times \mathbb{Z}_5$	6_{39}	$S^2 : (2, 1) (2, 1) (4, 1)$	$A_{4,-5}, A_{2,-3 4,-1}$	$\mathbb{Z}_2 \oplus \mathbb{Z}_{10}$

Δ	3-Manifold	Mat.	Seifert Structure	Triangulations	Homology
6	$S^3/Q_{16} \times \mathbb{Z}_7$	6 ₄₀	$S^2 : (2, 1) (2, 1) (4, 3)$	$A_{4,-7}, A_{2,-3 4,-3}$	$\mathbb{Z}_2 \oplus \mathbb{Z}_{14}$
	$S^3/Q_{20} \times \mathbb{Z}_7$	6 ₄₁	$S^2 : (2, 1) (2, 1) (5, 2)$	$A_{5,-7}, A_{2,-3 5,-2}$	\mathbb{Z}_{28}
	S^3/Q_{24}	6 ₄₂	$S^2 : (2, 1) (2, 1) (6, -5)$	\tilde{C}_6	$\mathbb{Z}_2 \oplus \mathbb{Z}_2$
	$S^3/Q_{28} \times \mathbb{Z}_3$	6 ₄₃	$S^2 : (2, 1) (2, 1) (7, -4)$	$A_{7,-3}, X_{2 3,-1}, J_{1 3,-4}$	\mathbb{Z}_{12}
	$S^3/Q_{28} \times \mathbb{Z}_5$	6 ₄₄	$S^2 : (2, 1) (2, 1) (7, -2)$	$A_{7,-5}, J_{1 5,-2}$	\mathbb{Z}_{20}
	$S^3/Q_{32} \times \mathbb{Z}_3$	6 ₄₅	$S^2 : (2, 1) (2, 1) (8, -5)$	$A_{8,-3}, X_{2 3,-2}, J_{1 3,-5}$	$\mathbb{Z}_2 \oplus \mathbb{Z}_6$
	$S^3/Q_{32} \times \mathbb{Z}_5$	6 ₄₆	$S^2 : (2, 1) (2, 1) (8, -3)$	$A_{8,-5}, J_{1 5,-3}$	$\mathbb{Z}_2 \oplus \mathbb{Z}_{10}$
	S^3/D_{56}	6 ₄₇	$S^2 : (2, 1) (2, 1) (7, -5)$	$A_{7,-2}, X_{2 2,-3}, X_{3 2,-1},$ $J_{1 2,-5}$	\mathbb{Z}_8
	S^3/D_{80}	6 ₄₈	$S^2 : (2, 1) (2, 1) (5, -1)$	$A_{5,-4}, J_{1 4,-1}$	\mathbb{Z}_{16}
	S^3/D_{96}	6 ₄₉	$S^2 : (2, 1) (2, 1) (3, 5)$	$A_{3,-8}, A_{2,-5 3,-2},$ $A_{2,-3 3,-5},$ $A_{2,-3 2,-3 3,-2}$	\mathbb{Z}_{32}
	S^3/D_{112}	6 ₅₀	$S^2 : (2, 1) (2, 1) (7, -3)$	$A_{7,-4}, J_{1 4,-3}$	\mathbb{Z}_{16}
	S^3/D_{160}	6 ₅₁	$S^2 : (2, 1) (2, 1) (5, 3)$	$A_{5,-8}, A_{2,-3 5,-3}$	\mathbb{Z}_{32}
	$S^3/P_{24} \times \mathbb{Z}_7$	6 ₅₂	$S^2 : (2, 1) (3, 1) (3, 1)$	$A_{3,-4 3,-1},$ $A_{2,-3 3,-1 3,-1}$	\mathbb{Z}_{21}
	$S^3/P_{24} \times \mathbb{Z}_{11}$	6 ₅₃	$S^2 : (2, 1) (3, 2) (3, 2)$	$A_{3,-5 3,-2},$ $A_{2,-3 3,-2 3,-2}$	\mathbb{Z}_{33}
	$S^3/P_{48} \times \mathbb{Z}_5$	6 ₅₄	$S^2 : (2, 1) (3, 2) (4, -3)$	$A_{3,-2 4,-1}, J_{3 2,-1}$	\mathbb{Z}_{10}
	$S^3/P_{48} \times \mathbb{Z}_7$	6 ₅₅	$S^2 : (2, 1) (3, 1) (4, -1)$	$A_{3,-1 4,-3}, J_{2 3,-1}$	\mathbb{Z}_{14}
	$S^3/P_{48} \times \mathbb{Z}_{11}$	6 ₅₆	$S^2 : (2, 1) (3, 2) (4, -1)$	$A_{3,-2 4,-3}$	\mathbb{Z}_{22}
	$S^3/P_{120} \times \mathbb{Z}_7$	6 ₅₇	$S^2 : (2, 1) (3, 1) (5, -3)$	$A_{3,-1 5,-2}, J_{2 2,-3}, O_{6,1}$	\mathbb{Z}_7
	$S^3/P_{120} \times \mathbb{Z}_{13}$	6 ₅₈	$S^2 : (2, 1) (3, 1) (5, -2)$	$A_{3,-1 5,-3}, J_{2 3,-2}$	\mathbb{Z}_{13}
	$S^3/P_{120} \times \mathbb{Z}_{17}$	6 ₅₉	$S^2 : (2, 1) (3, 2) (5, -3)$	$A_{3,-2 5,-2}$	\mathbb{Z}_{17}
	$S^3/P_{120} \times \mathbb{Z}_{23}$	6 ₆₀	$S^2 : (2, 1) (3, 2) (5, -2)$	$A_{3,-2 5,-3}$	\mathbb{Z}_{23}
	S^3/P'_{216}	6 ₆₁	$S^2 : (2, 1) (3, 1) (3, 2)$	$A_{3,-5 3,-1},$ $A_{3,-4 3,-2},$ $A_{2,-3 3,-2 3,-1}$	\mathbb{Z}_{27}
	General SFS	6 ₆₂	$S^2 : (3, 1) (3, 1) (3, -1)$	$A_{3,-2 3,-1 3,-1}$	$\mathbb{Z}_3 \oplus \mathbb{Z}_3$
	General SFS	6 ₆₃	$S^2 : (3, 1) (3, 2) (3, -1)$	$A_{3,-2 3,-2 3,-1}$	$\mathbb{Z}_3 \oplus \mathbb{Z}_6$
	General SFS	6 ₆₄	$S^2 : (3, 2) (3, 2) (3, -1)$	$A_{3,-2 3,-2 3,-2}$	$\mathbb{Z}_3 \oplus \mathbb{Z}_9$

Δ	3-Manifold	Mat.	Seifert Structure	Triangulations	Homology
6	$T^2 \times I / \begin{bmatrix} 1 & -1 \\ 1 & 0 \end{bmatrix}$	6 ₆₅	$S^2 : (2, 1) (3, 1) (6, -5)$	P_1	\mathbb{Z}
	$T^2 \times I / \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$	6 ₆₆	$S^2 : (2, 1) (4, 1) (4, -3)$	$C_{3,3}, T_{\text{II} 0,1 -1,0}$	$\mathbb{Z} \oplus \mathbb{Z}_2$
	$T^2 \times I / \begin{bmatrix} 0 & 1 \\ -1 & -1 \end{bmatrix}$	6 ₆₇	$S^2 : (3, 1) (3, 1) (3, -2)$	$A_{3,-1 3,-1 3,-1},$ $T_{\text{I} -1,1 -1,0}$	$\mathbb{Z} \oplus \mathbb{Z}_3$
	$T^2 \times I / \begin{bmatrix} -1 & 0 \\ -1 & -1 \end{bmatrix}$	6 ₆₈	$K^2 : (1, 1)$	$T_{\text{II} -1,0 1,-1},$ $Q_{K_0,K_0 -1,0 1,1}$	$\mathbb{Z} \oplus \mathbb{Z}_4$
	$T^2 \times I / \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$	6 ₆₉	$T^2 : (1, 1)$	$T_{\text{II} 1,-1 0,1}$	$\mathbb{Z} \oplus \mathbb{Z}$
	$T^2 \times I / \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}$	6 ₇₀	$S^2 : (2, 1) (2, 1) \dots$ $\dots (2, -1) (2, -1)$	$T_{\text{I} -1,0 0,-1},$ $Q_{K_0,K_0 1,0 0,1},$ $O_{6,2}$	$\mathbb{Z} \oplus \mathbb{Z}_2 \oplus \mathbb{Z}_2$
	$T^2 \times S^1$	6 ₇₁	$T^2 \times S^1$	$T_{\text{I} 1,0 0,1}$	$\mathbb{Z} \oplus \mathbb{Z} \oplus \mathbb{Z}$
	$(K^2 \times I)^2 / \begin{bmatrix} -1 & 0 \\ -1 & 1 \end{bmatrix}$	6 ₇₂	$S^2 : (2, 1) (2, 1) \dots$ $\dots (2, 1) (2, -1)$	$Q_{K_0,K_0 1,1 0,-1}$	$\mathbb{Z}_2 \oplus \mathbb{Z}_2 \oplus \mathbb{Z}_4$
	$(K^2 \times I)^2 / \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$	6 ₇₃	$\mathbb{R}P^2 : (2, 1), (2, -1)$	$Q_{K_0,K_0 0,-1 -1,0},$ $O_{6,3}, O_{6,4}, O_{6,5}$	$\mathbb{Z}_4 \oplus \mathbb{Z}_4$
	$(K^2 \times I)^2 / \begin{bmatrix} -1 & 1 \\ -1 & 0 \end{bmatrix}$	6 ₇₄	$\mathbb{R}P^2 : (2, 1), (2, 1)$	$Q_{K_0,K_0 0,1 -1,-1},$ $O_{6,6}$	$\mathbb{Z}_4 \oplus \mathbb{Z}_4$

Table A.2: Closed non-orientable prime minimal triangulations

Δ	3-Manifold	Triangulations	Homology
2	$S^2 \times S^1$	N_2	\mathbb{Z}
3	$\mathbb{R}P^2 \times S^1$	$N_{3,1}, N_{3,2}$	$\mathbb{Z} \oplus \mathbb{Z}_2$
6	$T^2 \times I / \begin{bmatrix} -1 & 1 \\ 1 & 0 \end{bmatrix}$	$T_{\text{II} -1,1 1,0}$	\mathbb{Z}
	$T^2 \times I / \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$	$(T_{\text{I} -1,0 -1,1} = K_{\text{I} 1,0 0,1}), T_{\text{I} 0,-1 -1,0},$ $(T_{\text{I} 0,1 1,0} = K_{\text{II} 0,-1 -1,0}), T_{\text{I} 1,0 1,-1},$ $K_{\text{III} 0,-1 -1,0}, Q_{T,T 0,1 -1,-1}$	$\mathbb{Z} \oplus \mathbb{Z}$
	$K^2 \times S^1$	$(T_{\text{II} 1,0 0,-1} = K_{\text{I} -1,0 -1,1}), K_{\text{II} 1,0 0,1}, K_{\text{III} 1,0 0,1}$	$\mathbb{Z} \oplus \mathbb{Z} \oplus \mathbb{Z}_2$
	$K^2 \times I / \begin{bmatrix} -1 & 1 \\ 0 & -1 \end{bmatrix}$	$K_{\text{I} 1,0 0,-1}, K_{\text{II} 0,1 1,0}, K_{\text{III} 0,1 1,0},$ $Q_{T,K_0 0,1 -1,-1}, Q_{T,K_0 1,0 0,1}, N_{6,1}, N_{6,2}, N_{6,3}, N_{6,4}$	$\mathbb{Z} \oplus \mathbb{Z}_4$
	$K^2 \times I / \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$	$K_{\text{I} -1,0 -1,-1}, K_{\text{II} -1,0 0,-1}, K_{\text{III} -1,0 0,-1},$ $Q_{K_{\text{II}},K_{\text{II}} 1,0 0,1}, Q_{T,K_0 -1,-1 1,0}$	$\mathbb{Z} \oplus \mathbb{Z}_2 \oplus \mathbb{Z}_2$

Appendix B

Tables of Normal Surfaces

Presented here are the vertex normal surfaces of every closed prime minimal triangulation formed from at most six tetrahedra. These vertex surfaces are enumerated first in the standard triangle-quadrilateral coordinate system described in Section 1.1 and then in the quadrilateral-only coordinate system of Tollefson [45].

The individual tables are arranged first by orientability and then by number of tetrahedra. Within each table triangulations are arranged according to method of construction. The individual triangulations are named as described in Sections 3.4 and 3.5.

For each surface several properties are listed. These include the Euler characteristic (marked χ in the tables), whether the surface is orientable and whether the surface is one-sided or two-sided. In the *Misc* column of each table, surfaces that are vertex links, edge links or splitting surfaces are marked as such. Note that in this context an edge link refers only to the boundary of a small neighbourhood of an edge, and not to the thick edge links described by Jaco and Rubinstein [17].

Table B.1: Normal surfaces in standard coordinates (1-tetrahedron orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{1,0}$	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{4,1}$	0	Non-orbl	1	Edge / Split	2	Orbl	2	Vertex
$L_{5,2}$	2	Orbl	2	Vertex				
C_1	0	Orbl	2	Edge / Split	2	Orbl	2	Vertex
	2	Orbl	2	Vertex				

Table B.2: Normal surfaces in standard coordinates (2-tetrahedron orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{0,1}$	0	Non-orbl	1	Edge / Split	2	Orbl	2	
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{2,1}$	-1	Non-orbl	1	Split	1	Non-orbl	1	
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L^{(1)}_{3,1}$	-2	Orbl	2		2	Orbl	2	
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L^{(2)}_{3,1}$	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{5,1}$	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{7,2}$	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{8,3}$	0	Non-orbl	1	Edge / Split	2	Orbl	2	Vertex
C_2	0	Orbl	2	Edge / Split	2	Orbl	2	Vertex
	1	Non-orbl	1	Edge / Split	2	Orbl	2	Vertex
	1	Non-orbl	1	Edge / Split				
\tilde{C}_2	0	Non-orbl	1	Edge / Split	0	Non-orbl	1	Edge / Split
	0	Non-orbl	1	Edge / Split	2	Orbl	2	Vertex
$L'(3,1)$	2	Orbl	2	Edge	2	Orbl	2	Vertex
	2	Orbl	2	Edge	2	Orbl	2	Vertex
	2	Orbl	2	Edge				

Table B.3: Normal surfaces in standard coordinates (3-tetrahedron orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{6,1}$	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{9,2}$	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{10,3}$	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{11,3}$	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{12,5}$	0	Non-orbl	1	Edge / Split	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{13,5}$	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
\tilde{C}_3	0	Non-orbl	1	Edge / Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				

Table B.4: Normal surfaces in standard coordinates (4-tetrahedron orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{7,1}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{11,2}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{13,3}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{14,3}$	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{15,4}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{16,7}$	0	Non-orbl	1	Edge / Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{17,5}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{18,5}$	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{19,7}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{21,8}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
\tilde{C}_4	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Non-orbl	1	Edge / Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$C_{2,2}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$A_{2,-3}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Non-orbl	1					
$A_{3,-2}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Orbl	2					
$J_{1 -2,1}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex

Table B.5: Normal surfaces in standard coordinates (5-tetrahedron orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{8,1}$	-2	Non-orbl	1	Split	0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	-2	Orbl	2					
$L_{13,2}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{16,3}$	-2	Non-orbl	1	Split	0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
$L_{17,3}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{17,4}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{19,4}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{20,9}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1	Edge / Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{22,5}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{23,5}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{23,7}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{24,7}$	-2	Non-orbl	1	Split	0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{25,7}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{25,9}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{26,7}$	-2	Orbl	2		0	Orbl	2	Edge
	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{27,8}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{29,8}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{29,12}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{30,11}$	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{31,12}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{34,13}$	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
\tilde{C}_5	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1	Edge / Split	2	Orbl	2	Vertex
$C_{2,3}$	-2	Orbl	2		0	Orbl	2	Edge
	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{2,-5}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
$A_{3,-5}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{3,-4}$	-2	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Orbl	2					
$A_{4,-3}$	-1	Non-orbl	1		0	Orbl	2	
	-1	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
$A_{5,-3}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{5,-2}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{2,-3 2,-3}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
$A_{2,-3 3,-2}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{2,-3 3,-1}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{3,-2 3,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{3,-2 3,-1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$X_{2 2,-1}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$J_{1 2,-3}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$J_{1 3,-2}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$J_{1 3,-1}$	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$J_{2 2,-1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
P_0	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
P'_0	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex

Table B.6: Normal surfaces in standard coordinates (6-tetrahedron orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{9,1}$	-4	Orbl	2		-2	Orbl	2	
	-4	Orbl	2		-2	Orbl	2	
	-4	Orbl	2		-2	Orbl	2	
	-4	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	-2	Orbl	2					
$L_{15,2}$	-4	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	-2	Orbl	2					
$L_{19,3}$	-4	Orbl	2		-2	Orbl	2	
	-4	Orbl	2		-2	Orbl	2	
	-4	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{20,3}$	-2	Non-orbl	1	Split	0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	-2	Orbl	2					
$L_{21,4}$	-4	Orbl	2		-2	Orbl	2	
	-4	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	-2	Orbl	2					
$L_{23,4}$	-4	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	-2	Orbl	2					
$L_{24,5}$	-2	Non-orbl	1	Split	0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	-2	Orbl	2					

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{24,11}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1	Edge / Split	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{27,5}$	-4	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	-2	Orbl	2					
$L_{28,5}$	-2	Non-orbl	1	Split	0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	-2	Orbl	2					
$L_{29,9}$	-4	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{30,7}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-1	Non-orbl	1	Split	2	Orbl	2	Vertex
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{31,7}$	-4	Orbl	2		-2	Orbl	2	
	-4	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
$L_{31,11}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{32,7}$	-2	Non-orbl	1	Split	0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	-2	Orbl	2					
$L_{33,7}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{33,10}$	-4	Orbl	2		0	Orbl	2	Edge
	-4	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	-2	Orbl	2					
$L_{34,9}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{35,8}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{36,11}$	-2	Non-orbl	1	Split	0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{37,8}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{37,10}$	-4	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
$L_{39,14}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{39,16}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{40,11}$	-2	Non-orbl	1	Split	0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
$L_{41,11}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{41,12}$	-4	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{41,16}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{43,12}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{44,13}$	-2	Non-orbl	1	Split	0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{45,19}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{46,17}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{47,13}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$L_{49,18}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{50,19}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$L_{55,21}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
\tilde{C}_6	-2	Non-orbl	1	Split	-2	Orbl	2	
	-2	Non-orbl	1	Split	0	Non-orbl	1	Edge / Split
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		2	Orbl	2	Vertex
$C_{3,3}$	-2	Non-orbl	1	Split	0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$A_{2,-7}$	-2	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
$A_{3,-8}$	-2	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{3,-7}$	-2	Non-orbl	1		0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Orbl	2					

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{4,-7}$	-2	Non-orbl	1		0	Orbl	2	
	-1	Non-orbl	1		0	Orbl	2	
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
$A_{4,-5}$	-2	Non-orbl	1		0	Orbl	2	
	-2	Non-orbl	1		0	Orbl	2	
	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Non-orbl	1					
$A_{5,-8}$	-2	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{5,-7}$	-2	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{5,-4}$	-2	Non-orbl	1		0	Non-orbl	1	
	-2	Orbl	2		0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
$A_{7,-5}$	-2	Orbl	2		0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Orbl	2					

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{7,-4}$	-2	Non-orbl	1		0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
$A_{7,-3}$	-2	Orbl	2		0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
$A_{7,-2}$	-2	Orbl	2		0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Orbl	2					
$A_{8,-5}$	-1	Non-orbl	1		0	Orbl	2	
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Orbl	2					
$A_{8,-3}$	-1	Non-orbl	1		0	Orbl	2	
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Orbl	2					
$A_{2,-5 2,-3}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{2,-5 3,-2}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$A_{2,-5 3,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{2,-3 3,-5}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$A_{2,-3 3,-4}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2					
$A_{2,-3 4,-3}$	-1	Non-orbl	1		0	Orbl	2	
	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{2,-3 4,-1}$	-1	Non-orbl	1		0	Orbl	2	
	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{2,-3 5,-3}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{2,-3 5,-2}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$A_{3,-5 3,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$A_{3,-5 3,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{3,-4 3,-2}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2					
$A_{3,-4 3,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2					
$A_{3,-2 4,-3}$	-2	Orbl	2		0	Orbl	2	
	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{3,-2 4,-1}$	-2	Orbl	2		0	Orbl	2	
	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{3,-2 5,-3}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$A_{3,-2 5,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$A_{3,-1 4,-3}$	-2	Orbl	2		0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{3,-1 5,-3}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{3,-1 5,-2}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{2,-3 2,-3 2,-3}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{2,-3 2,-3 3,-2}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$A_{2,-3 2,-3 3,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{2,-3 3,-2 3,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$A_{2,-3 3,-2 3,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2					
$A_{2,-3 3,-1 3,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{3,-2 3,-2 3,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$A_{3,-2 3,-2 3,-1}$	-1	Non-orbl	1		0	Orbl	2	
	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{3,-2 3,-1 3,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$A_{3,-1 3,-1 3,-1}$	-2	Non-orbl	1	Split	0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		2	Orbl	2	Vertex
$X_{2 2,-3}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$X_{2 3,-2}$	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$X_{2 3,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$X_{3 2,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$J_{1 2,-5}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$J_{1 3,-5}$	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$J_{1 3,-4}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
$J_{1 4,-3}$	-2	Non-orbl	1		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$J_{1 4,-1}$	-2	Non-orbl	1		0	Non-orbl	1	
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
$J_{1 5,-3}$	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$J_{1 5,-2}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$J_{2 2,-3}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$J_{2 3,-2}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$J_{2 3,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$J_{3 2,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
P_1	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$T_{I -1,0 0,-1}$	-2	Non-orbl	1	Split	0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Non-orbl	1					
$T_{I -1,1 -1,0}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$T_{I 1,0 0,1}$	-2	Non-orbl	1	Split	0	Orbl	2	
	0	Orbl	2		0	Orbl	2	
	0	Orbl	2		0	Orbl	2	
	0	Orbl	2		0	Orbl	2	
	0	Orbl	2		2	Orbl	2	Vertex
	0	Orbl	2					
$T_{II -1,0 1,-1}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$T_{II 0,1 -1,0}$	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2	Edge				
$T_{II 1,-1 0,1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$Q_{K_0,K_0 -1,0 1,1}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	
	0	Orbl	2		0	Orbl	2	
	0	Orbl	2		2	Orbl	2	Vertex
$Q_{K_0,K_0 0,-1 -1,0}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$Q_{K_0, K_0 0,1 -1,-1}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		2	Orbl	2	Vertex
$Q_{K_0, K_0 1,0 0,1}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Non-orbl	1					
$Q_{K_0, K_0 1,1 0,-1}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Non-orbl	1					
$O_{6,1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	2	Orbl	2	Vertex
$O_{6,2}$	-2	Non-orbl	1	Split	0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Non-orbl	1					
$O_{6,3}$	-2	Non-orbl	1	Split	0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		2	Orbl	2	Vertex

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$O_{6,4}$	-2	Non-orbl	1	Split	0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Non-orbl	1					
$O_{6,5}$	-2	Orbl	2		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Non-orbl	1					
$O_{6,6}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex

Table B.7: Normal surfaces in standard coordinates (2-tetrahedron non-orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
N_2	0	Non-orbl	2	Edge	2	Orbl	2	
	0	Orbl	1	Edge / Split	2	Orbl	2	Vertex

Table B.8: Normal surfaces in standard coordinates (3-tetrahedron non-orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$N_{3,1}$	-1	Non-orbl	1	Split	0	Orbl	1	Edge / Split
	0	Non-orbl	1	Edge / Split	1	Non-orbl	2	
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex
$N_{3,2}$	-1	Non-orbl	1	Split	0	Orbl	2	Edge
	0	Non-orbl	1		1	Non-orbl	2	
	0	Orbl	1		2	Orbl	2	Vertex

Table B.9: Normal surfaces in standard coordinates (6-tetrahedron non-orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$T_{I -1,0 -1,1}$	-2	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	1		2	Orbl	2	Vertex
	0	Orbl	1					
$T_{I 0,-1 -1,0}$	-2	Non-orbl	1	Splitting	0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Orbl	1		2	Orbl	2	Vertex
	0	Orbl	1					
$T_{I 0,1 1,0}$	-2	Non-orbl	1	Splitting	0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Orbl	1		2	Orbl	2	Vertex
	0	Orbl	1					
$T_{I 1,0 1,-1}$	-2	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex
	0	Non-orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$T_{\Pi -1,1 1,0}$	0	Non-orbl	2	Edge	0	Orbl	2	Edge
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex
	0	Orbl	2					
$T_{\Pi 1,0 0,-1}$	-1	Non-orbl	1		0	Non-orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex
	0	Non-orbl	2	Edge				
$K_{I -1,0 -1,-1}$	-1	Non-orbl	1		0	Non-orbl	2	Edge
	-1	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	2		2	Orbl	2	Vertex
	0	Non-orbl	2	Edge				
$K_{I 1,0 0,-1}$	-2	Non-orbl	2		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex
	0	Non-orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$K_{II -1,0 0,-1}$	-2	Non-orbl	1	Splitting	0	Non-orbl	2	
	-1	Non-orbl	1	Splitting	0	Non-orbl	2	Edge
	-1	Non-orbl	1	Splitting	0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
$K_{II 0,1 1,0}$	-2	Orbl	1	Splitting	0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		2	Orbl	2	Vertex
	0	Non-orbl	2	Edge				
$K_{II 1,0 0,1}$	-2	Orbl	1	Splitting	0	Non-orbl	2	Edge
	-1	Non-orbl	1	Splitting	0	Non-orbl	2	Edge
	-1	Non-orbl	1	Splitting	0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2	Edge	0	Orbl	2	
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex
$K_{III -1,0 0,-1}$	-2	Orbl	1	Splitting	0	Non-orbl	2	
	-2	Orbl	2		0	Non-orbl	2	
	-2	Orbl	2		0	Non-orbl	2	
	-1	Non-orbl	1		0	Orbl	1	
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Non-orbl	1					

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$K_{III 0,-1 -1,0}$	-2	Orbl	1	Splitting	0	Orbl	1	
	-2	Orbl	2		0	Orbl	1	
	-2	Orbl	2		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Non-orbl	2		2	Orbl	2	Vertex
	0	Orbl	1					
$K_{III 0,1 1,0}$	-2	Non-orbl	1	Splitting	0	Non-orbl	2	Edge
	-2	Non-orbl	2		0	Non-orbl	2	Edge
	-2	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		2	Orbl	2	Vertex
	0	Non-orbl	2					
$K_{III 1,0 0,1}$	-2	Non-orbl	1	Splitting	0	Non-orbl	2	Edge
	-2	Non-orbl	2		0	Non-orbl	2	Edge
	-2	Non-orbl	2		0	Orbl	1	
	-1	Non-orbl	1		0	Orbl	1	
	-1	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Non-orbl	2		2	Orbl	2	Vertex
	0	Non-orbl	2					
$Q_{T,T 0,1 -1,-1}$	-2	Non-orbl	1		0	Non-orbl	2	Edge
	-2	Non-orbl	1		0	Orbl	1	
	-2	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2		2	Orbl	2	Vertex
	0	Non-orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$Q_{T,K_0 -1,-1 1,0}$	-1	Non-orbl	1		0	Non-orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	2		2	Orbl	2	Vertex
	0	Non-orbl	2		2	Orbl	2	Vertex
$Q_{T,K_0 0,1 -1,-1}$	-2	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex
$Q_{T,K_0 1,0 0,1}$	-2	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex
$Q_{K_{II},K_{II} 1,0 0,1}$	-2	Non-orbl	1	Splitting	0	Non-orbl	1	
	-2	Non-orbl	2		0	Non-orbl	2	
	-2	Non-orbl	2		0	Non-orbl	2	Edge
	-1	Non-orbl	1		0	Non-orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		2	Orbl	2	Vertex
$N_{6,1}$	-2	Non-orbl	1	Splitting	0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		2	Orbl	2	Vertex
	0	Non-orbl	2	Edge				
	0	Non-orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$N_{6,2}$	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	1	
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex
$N_{6,3}$	-2	Non-orbl	1	Splitting	0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex
$N_{6,4}$	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	1	
	0	Non-orbl	2	Edge	2	Orbl	2	Vertex

Table B.10: Normal surfaces in quadrilateral coordinates (1-tetrahedron orientable)

Triang.	χ	Orient.	Sides	Misc.
$L_{1,0}$	0	Orbl	2	Edge
$L_{4,1}$	0	Non-orbl	1	Edge / Splitting
C_1	0	Orbl	2	Edge / Splitting

Table B.11: Normal surfaces in quadrilateral coordinates (2-tetrahedron orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{0,1}$	0	Non-orbl	1	Edge / Splitting	2	Orbl	2	
	0	Orbl	2	Edge				
$L_{2,1}$	0	Orbl	2	Edge	1	Non-orbl	1	
$L^{(1)}_{3,1}$	0	Orbl	2	Edge	2	Orbl	2	
$L^{(2)}_{3,1}$	0	Orbl	2	Edge				
$L_{5,1}$	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{7,2}$	0	Orbl	2	Edge				
$L_{8,3}$	0	Non-orbl	1	Edge / Splitting				
C_2	0	Orbl	2	Edge / Splitting	1	Non-orbl	1	Edge / Splitting
	1	Non-orbl	1	Edge / Splitting				
\tilde{C}_2	0	Non-orbl	1	Edge / Splitting	0	Non-orbl	1	Edge / Splitting
	0	Non-orbl	1	Edge / Splitting				
$L'(3,1)$	2	Orbl	2	Edge	2	Orbl	2	Edge
	2	Orbl	2	Edge				

Table B.12: Normal surfaces in quadrilateral coordinates (3-tetrahedron orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{6,1}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{9,2}$	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{10,3}$	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{11,3}$	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{12,5}$	0	Non-orbl	1	Edge / Splitting	0	Orbl	2	Edge
$L_{13,5}$	0	Orbl	2	Edge	0	Orbl	2	Edge
\tilde{C}_3	0	Non-orbl	1	Edge / Splitting	0	Orbl	2	Edge
	0	Orbl	2	Edge				

Table B.13: Normal surfaces in quadrilateral coordinates (4-tetrahedron orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{7,1}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{11,2}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{13,3}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{14,3}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{15,4}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{16,7}$	0	Non-orbl	1	Edge / Splitting	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{17,5}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{18,5}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{19,7}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{21,8}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
\tilde{C}_4	0	Non-orbl	1	Edge / Splitting	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$C_{2,2}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$A_{2,-3}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
$A_{3,-2}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$J_{1 -2,1}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge				

Table B.14: Normal surfaces in quadrilateral coordinates (5-tetrahedron orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{8,1}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{13,2}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{16,3}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{17,3}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{17,4}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{19,4}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{20,9}$	0	Non-orbl	1	Edge / Splitting	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{22,5}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{23,5}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{23,7}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{24,7}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{25,7}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{25,9}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{26,7}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{27,8}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{29,8}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{29,12}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{30,11}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{31,12}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{34,13}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
\tilde{C}_5	0	Non-orbl	1	Edge / Splitting	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$C_{2,3}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$A_{2,-5}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1					
$A_{3,-5}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2					
$A_{3,-4}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2					
$A_{4,-3}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2					

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{5,-3}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2					
$A_{5,-2}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2					
$A_{2,-3 2,-3}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1					
$A_{2,-3 3,-2}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2					
$A_{2,-3 3,-1}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2					
$A_{3,-2 3,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2					
$A_{3,-2 3,-1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2					
$X_{2 2,-1}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$J_{1 2,-3}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$J_{1 3,-2}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$J_{1 3,-1}$	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge				
$J_{2 2,-1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
P_0	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
P'_0	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				

Table B.15: Normal surfaces in quadrilateral coordinates (6-tetrahedron orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{9,1}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$L_{15,2}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{19,3}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{20,3}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{21,4}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{23,4}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{24,5}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{24,11}$	0	Non-orbl	1	Edge / Splitting	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{27,5}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{28,5}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{29,9}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{30,7}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{31,7}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{31,11}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{32,7}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{33,7}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{33,10}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{34,9}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{35,8}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{36,11}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{37,8}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{37,10}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{39,14}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{39,16}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{40,11}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{41,11}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{41,12}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{41,16}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{43,12}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{44,13}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$L_{45,19}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{46,17}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{47,13}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{49,18}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{50,19}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$L_{55,21}$	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
\tilde{C}_6	0	Non-orbl	1	Edge / Splitting	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$C_{3,3}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$A_{2,-7}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{3,-8}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-7}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{4,-7}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{4,-5}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{5,-8}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{5,-7}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{5,-4}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{7,-5}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{7,-4}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{7,-3}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{7,-2}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{8,-5}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{8,-3}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{2,-5 2,-3}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
$A_{2,-5 3,-2}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{2,-5 3,-1}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{2,-3 3,-5}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{2,-3 3,-4}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{2,-3 4,-3}$	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2					

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{2,-3 4,-1}$	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2					
$A_{2,-3 5,-3}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{2,-3 5,-2}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-5 3,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-5 3,-1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-4 3,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-4 3,-1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-2 4,-3}$	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2					
$A_{3,-2 4,-1}$	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2					

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{3,-2 5,-3}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-2 5,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-1 4,-3}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-1 5,-3}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-1 5,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{2,-3 2,-3 2,-3}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
$A_{2,-3 2,-3 3,-2}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{2,-3 2,-3 3,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2					
$A_{2,-3 3,-2 3,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$A_{2,-3 3,-2 3,-1}$	-2	Orbl	2		0	Orbl	2	
	-2	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{2,-3 3,-1 3,-1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-2 3,-2 3,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-2 3,-2 3,-1}$	-1	Non-orbl	1		0	Orbl	2	
	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2					
$A_{3,-2 3,-1 3,-1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
$A_{3,-1 3,-1 3,-1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2					
$X_{2 2,-3}$	-2	Orbl	2		0	Orbl	2	Edge
	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$X_{2 3,-2}$	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge				
$X_{2 3,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$X_{3 2,-1}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$J_{1 2,-5}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$J_{1 3,-5}$	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$J_{1 3,-4}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$J_{1 4,-3}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$J_{1 4,-1}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$J_{1 5,-3}$	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$J_{1 5,-2}$	-2	Orbl	2		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$J_{2 2,-3}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$J_{2 3,-2}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$J_{2 3,-1}$	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$J_{3 2,-1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
P_1	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$T_{I -1,0 0,-1}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1					
$T_{I -1,1 -1,0}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
$T_{I 1,0 0,1}$	0	Orbl	2		0	Orbl	2	
	0	Orbl	2		0	Orbl	2	
	0	Orbl	2		0	Orbl	2	
	0	Orbl	2		0	Orbl	2	
	0	Orbl	2					
$T_{II -1,0 1,-1}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$T_{II 0,1 -1,0}$	-1	Non-orbl	1		0	Orbl	2	Edge
	-1	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$T_{II 1,-1 0,1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				
$Q_{K_0,K_0 -1,0 1,1}$	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	
	0	Orbl	2		0	Orbl	2	
	0	Orbl	2					
$Q_{K_0,K_0 0,-1 -1,0}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1					
$Q_{K_0,K_0 0,1 -1,-1}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1					
$Q_{K_0,K_0 1,0 0,1}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1		0	Orbl	2	Edge
$Q_{K_0,K_0 1,1 0,-1}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
$O_{6,1}$	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2		0	Orbl	2	Edge
	0	Orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$O_{6,2}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	1					
$O_{6,3}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
$O_{6,4}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1					
$O_{6,5}$	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Non-orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1					
$O_{6,6}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1					

Table B.16: Normal surfaces in quadrilateral coordinates (2-tetrahedron non-orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
N_2	0	Non-orbl	2	Edge	2	Orbl	2	
	0	Orbl	1	Edge / Split				

Table B.17: Normal surfaces in quadrilateral coordinates (3-tetrahedron non-orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$N_{3,1}$	0	Non-orbl	1	Edge / Split	0	Orbl	1	Edge / Split
	0	Non-orbl	2	Edge	1	Non-orbl	2	
$N_{3,2}$	0	Non-orbl	1		0	Orbl	2	Edge
	0	Orbl	1		1	Non-orbl	2	

Table B.18: Normal surfaces in quadrilateral coordinates (6-tetrahedron non-orientable)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$T_{I -1,0 -1,1}$	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	2	
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
	0	Orbl	1		0	Orbl	2	Edge
	0	Orbl	1					
$T_{I 0,-1 -1,0}$	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Orbl	1		0	Orbl	2	Edge
	0	Orbl	1					
$T_{I 0,1 1,0}$	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Orbl	1		0	Orbl	2	Edge
	0	Orbl	1					
$T_{I 1,0 1,-1}$	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
	0	Non-orbl	2	Edge				
$T_{II -1,1 1,0}$	0	Non-orbl	2	Edge	0	Orbl	2	
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
	0	Non-orbl	2	Edge	0	Orbl	2	Edge

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$T_{II 1,0 0,-1}$	-1	Non-orbl	1		0	Non-orbl	2	Edge
	-1	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	2	
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
$K_{I -1,0 -1,-1}$	-1	Non-orbl	1		0	Non-orbl	2	Edge
	-1	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	2		0	Orbl	2	Edge
$K_{I 1,0 0,-1}$	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
	0	Non-orbl	2	Edge				
$K_{II -1,0 0,-1}$	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	2		0	Orbl	2	Edge
$K_{II 0,1 1,0}$	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$K_{II 1,0 0,1}$	0	Non-orbl	2		0	Non-orbl	2	Edge
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	2	
	0	Non-orbl	2	Edge	0	Orbl	2	
$K_{III -1,0 0,-1}$	0	Non-orbl	1		0	Non-orbl	2	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	1		0	Orbl	2	Edge
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Non-orbl	2					
$K_{III 0,-1 -1,0}$	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	Edge
	0	Orbl	1		0	Orbl	2	Edge
	0	Orbl	1					
$K_{III 0,1 1,0}$	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge				
$K_{III 1,0 0,1}$	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
	0	Non-orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$Q_{T,T 0,1 -1,-1}$	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	2	
	0	Non-orbl	2	Edge				
$Q_{T,K_0 -1,-1 1,0}$	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	
	0	Non-orbl	2	Edge				
$Q_{T,K_0 0,1 -1,-1}$	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	1	
$Q_{T,K_0 1,0 0,1}$	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	2	Edge
$Q_{K_{II},K_{II} 1,0 0,1}$	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	2	
	0	Non-orbl	2		0	Orbl	2	
$N_{6,1}$	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge				

(continued on next page)

(continued from previous page)

Triang.	χ	Orient.	Sides	Misc.	χ	Orient.	Sides	Misc.
$N_{6,2}$	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	1	
	0	Non-orbl	2	Edge				
$N_{6,3}$	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	1	
$N_{6,4}$	0	Non-orbl	1		0	Non-orbl	2	Edge
	0	Non-orbl	1		0	Orbl	1	
	0	Non-orbl	2		0	Orbl	1	
	0	Non-orbl	2	Edge	0	Orbl	1	
	0	Non-orbl	2	Edge				