

DISCRETE MORSE THEORY ON MODULI SPACES OF PLANAR POLYGONS

A thesis submitted to

Chennai Mathematical Institute

in partial fulfilment of the requirements for the degree of

MASTER OF SCIENCE

in

MATHEMATICS

by

Nachiketa Adhikari

SUPERVISOR: Priyavrat Deshpande

April 2017

Abstract

Given n positive real numbers l_1, \dots, l_n , the tuple

$$L = (l_1, \dots, l_n) \in \mathbb{R}_+^n$$

is called a *length vector*. The set of all planar polygons with these numbers as side lengths has a natural structure of a topological space. If we consider two polygons to be equivalent when one can be obtained from the other by a rotation or translation of the plane, the resulting set of equivalence classes of polygons can be the empty set, a manifold of dimension $n - 3$, or a manifold of dimension $n - 3$ with finitely many singularities. This space is called the moduli space of polygons associated with the given length vector L , and is denoted by \mathcal{M}_L .

Moduli spaces of polygons have been encountered in engineering contexts for centuries. Robot arms, for example, can be profitably modeled as polygon spaces. Many of the questions asked in the topological theory have direct applications in robotics. An important result in this context is that of M. Farber and D. Schütz, which states that the \mathbb{Z} -homology groups of \mathcal{M}_L are all free, and even gives a formula for the ranks of these homology groups purely in terms of L .

The aim of this thesis is to understand a natural regular cell structure admitted by these moduli spaces when they are manifolds. This cell structure, introduced by G. Panina, has a convenient combinatorial description which can be exploited to extract information about the space. In fact, the k -cells correspond precisely to cyclically ordered partitions of the set $\{1, \dots, n\}$ into $n - 3 - k$ blocks. (For example, when $n = 4$, the partition $\{1, 2\}\{3\}\{4\}$ corresponds to a 0-cell, which is the same as the partition $\{3\}\{4\}\{1, 2\}$, which is the same as the partition $\{4\}\{1, 2\}\{3\}$.)

After a brief glance at discrete Morse theory, we study Panina's construction of a so-called "perfect" Morse function on this cell structure. This perfect Morse function pairs off cells in an orderly fashion, until precisely as many k -cells remain, for each k , as the k^{th} Betti number of \mathcal{M}_L . On the way, we also use Panina's results to gain combinatorial insights into why Farber's and Schütz's result holds.

The moduli spaces we are studying admit a natural involution, wherein each polygon is mapped to its reflection about the X -axis. Generically, this is fixed-point free, and is also otherwise sufficiently nice, so that the quotient by \mathbb{Z}_2 is a manifold. We show that Panina's cell structure interacts comfortably with this action to give a cell structure on the quotient. We prove, in certain special cases, the existence of \mathbb{Z}_2 -perfect Morse functions on these quotient manifolds, using results by J.-C. Hausmann and A. Knutson.

Acknowledgments

I am deeply indebted to my adviser, Prof. Priyavrat Deshpande, for his support and guidance at every step. His infectious enthusiasm, attention to detail and hands-on approach ensured that this was a mathematically enriching and enjoyable experience throughout.

I thank my teachers at CMI, who ensured I was always in a stimulating and challenging learning environment. I mention here especially Prof. T. R. Ramadas and Prof. Vijay Ravikumar, whose helpful advice I have greatly benefited from. Thanks are also due to my fellow students, with whom I had many fruitful discussions.

My family has always encouraged and stood by me in all my endeavors, and this thesis would not have been possible without them.

Contents

Abstract	i
Acknowledgments	ii
Introduction	iv
1 Moduli spaces of planar polygons	1
1.1 Preliminaries	1
1.2 Basic results	4
2 A regular cell structure	8
2.1 Labeling polygons	8
2.2 The cells	13
2.3 Permutohedra	18
2.4 Face posets and dual complexes	19
3 Discrete Morse theory	22
3.1 A primer on discrete Morse theory	22
3.2 A perfect Morse function on \mathcal{K}_L^*	29
4 The \mathbb{Z}_2-action on \mathcal{M}_L	37
4.1 \mathbb{Z}_2 -equivariance	37
4.2 The \mathbb{Z}_2 -homology of \mathcal{O}_L	39
4.3 Discrete Morse theory on \mathcal{C}_L^*	42
5 Looking ahead	49
5.1 A minimal complex for \mathcal{M}_L	49
5.2 Topology of \mathcal{O}_L	51
5.3 Other directions	52
Bibliography	54

Introduction

The study of moduli spaces of polygons has its origins in the mechanical linkages used in engineering. A *mechanical linkage* is an assembly of bodies connected to manage forces and movement. A lever, for example, is a simple mechanical linkage. Topologically, the natural way to think about such linkages is as (possibly self-intersecting) polygons with some constraints. Generally speaking, the *moduli space* of a system is the space of all possible states of that system. Thus the space of all possible configurations of a mechanical linkage is naturally identified with a moduli space of polygons.

The topological study of these moduli spaces was initiated in the 1960s. Over the course of half a century, many techniques—including Morse theory and symplectic geometry, among others—have been used to understand these spaces and compute their Betti numbers as well as other topological invariants. Over time, many of these results have found applications in molecular biology, statistical shape analysis and, most notably, robotics: it turns out that the study of robot arms, which are themselves mechanical linkages, can benefit greatly from the topological theory. Consequently, the study of such moduli spaces can be considered part of the blossoming discipline called *topological robotics*, elaborated upon by M. Farber in [Far08].

The study of polygons in Euclidean space \mathbb{R}^3 has a rich history, with many interesting results. However, we restrict our attention to moduli spaces of polygons in the plane, that is, \mathbb{R}^2 . Since we are concerned only with the shapes of the polygons we get, and not with how they lie within the ambient space, we usually consider these moduli spaces modulo rotations and translations of the plane. In such a situation the moduli space, denoted by \mathcal{M}_L , is, generically, a manifold. For example, modulo rotations and translations of the plane, there are only two triangles with side lengths 3, 4 and 5, so the corresponding moduli space consists of two points. In other words, \mathcal{M}_L (where L depends on the lengths 3, 4 and 5) is a 0-sphere.

Computations of Betti numbers and other topological invariants of \mathcal{M}_L are among the classical results of the field, while in [Pan12], G. Panina describes a regular CW complex structure on \mathcal{M}_L (a *regular CW complex* is one in which the attaching maps are all homeomorphisms). The cells and boundary maps of this CW complex have combinatorial descriptions which lend themselves to easy comprehension and computation, and, moreover, provide a direct connection between the polygons in \mathcal{M}_L and the cells they belong to, allowing one to “see” how one can move from a certain cell to the other by adjusting the polygon accordingly.

An analog of classical Morse theory, known as *discrete Morse theory*, can be used to eliminate cells in a regular CW complex that do not contribute to homology. A discrete Morse function, as described by R. Forman in [For02], is a collection of pairs of cells such that each cell appears

in at most one pair. The cells that do not appear in any pair are the ones that might possibly contribute to homology. When the number of unpaired cells is equal to the sum of Betti numbers of the complex, we say that the function is *perfect*. Discrete Morse functions are useful because they allow us to construct complexes with fewer cells while maintaining the homotopy type.

The space \mathcal{M}_L admits an involution—reflection in the X -axis—that, generically, has no fixed points. Therefore we can form the quotient $\mathcal{O}_L = \mathcal{M}_L / \mathbb{Z}_2$. Adapting the description in [Pan12], we construct a CW complex on \mathcal{O}_L and show that the cells have a convenient combinatorial description, just as the ones in \mathcal{M}_L do. In [PZ15], G. Panina describes a perfect Morse function on the dual CW complex of \mathcal{M}_L . We modify this function to construct a discrete Morse function on the dual CW complex of \mathcal{O}_L , and show that in some cases the function is perfect.

Chapter-wise organization

Chapter 1. The first chapter contains basic definitions needed to understand most of the results. We also introduce here some of the foundational, well-known results of the area (taken mainly from [Far08] and [KM95]). Examples of familiar manifolds arising as \mathcal{M}_L , such as spheres, tori and orientable surfaces, are discussed.

Chapter 2. With this chapter, the main content of the thesis begins. We discuss in detail the construction, following [Pan12], of a regular CW complex on \mathcal{M}_L . Included here is a new proof of the fact that the convex polygons in \mathcal{M}_L form an open ball. Combinatorial aspects of the CW complex are illuminated using permutohedra, face figures and dual posets.

Chapter 3. We start the third chapter with a brief tour of discrete Morse theory as done in [For02], taking care to formulate the main results in the language of posets as well as that of directed graphs. Then we examine the perfect Morse function outlined in [PZ15]. A standard example (that of the circle) is used throughout Chapters 2 and 3 to illustrate the ideas.

Chapter 4. The quotient space \mathcal{O}_L is the focus of this chapter. We describe a CW complex structure on \mathcal{O}_L : this in fact follows from the fact that the CW structure on \mathcal{M}_L is \mathbb{Z}_2 -equivariant. Then, as in Chapter 3, we construct a discrete Morse function on the dual complex, and look at some cases where this function is actually perfect (with respect to \mathbb{Z}_2 -Betti numbers). We also give a necessary and sufficient condition for this function to arise from a \mathbb{Z}_2 -equivariant discrete Morse function on (the dual complex of) \mathcal{M}_L . With the exception of a calculation of \mathbb{Z}_2 -Betti numbers of \mathcal{O}_L based on [HK98], all results in this chapter are, to the best of our knowledge, new.

Chapter 5. Finally, we describe some of our current efforts and outline avenues for possible future work. Some of these are natural extensions of the questions answered in the first four chapters. Others move in slightly different directions.

Chapter 1

Moduli spaces of planar polygons

In this chapter, we introduce the basic definitions and examples. We also state some of the foundational results of the area. The main source used is M. Farber's [Far08].

Notation: The set $\{1, \dots, n\}$ will be denoted throughout by $[n]$.

1.1 Preliminaries

Before we move to the study of moduli spaces, we need a precise definition of what we mean by the word *polygon*.

Definition 1.1. A **length vector**

$$L = (l_1, \dots, l_n) \in \mathbb{R}_+^n$$

is an n -tuple of positive real numbers.

Definition 1.2. A **polygon with length vector** L ,

$$P = (\vec{u}_1, \dots, \vec{u}_n) \in \mathbb{R}^2 \times \dots \times \mathbb{R}^2, \quad \|\vec{u}_i\| = 1 \quad \forall i,$$

is an n -tuple of unit-length vectors in the plane, satisfying the equation

$$\sum_{i=1}^n l_i \vec{u}_i = 0. \tag{1.1}$$

The points

$$\begin{aligned} V_1 &= 0, \\ V_2 &= l_1 \vec{u}_1, \\ V_3 &= l_1 \vec{u}_1 + l_2 \vec{u}_2, \\ &\vdots \\ \text{and } V_n &= \sum_{i=1}^{n-1} l_i \vec{u}_i \end{aligned}$$

are called the **vertices** of the polygon, while the segments joining successive edges (thought of as translated vectors) are called the **sides** or **edges** of the polygon.

Thus a polygon always has its first vertex at (and therefore last edge pointing towards) the origin. Observe that our edges are always directed.

The way we have defined them, polygons can be objects which we do not conventionally recognize as polygons, as the following example and figures show.

Example 1.1. Let $L = (1, 2, 1, 2)$. Fig. 1.1 shows examples of polygons with this length vector. Fig. 1.1a shows an object the like of which we usually associate with the word “polygon”. However, polygons according to our definition can be non-convex and even self-intersecting, as Fig. 1.1b shows. In fact, our polygons can even degenerate to a line (Fig. 1.1c).




Figure 1.1: Examples of polygons with length vector $L = (1, 2, 1, 2)$.

Example 1.2. Here we list some non-examples. Again, let $L = (1, 2, 1, 2)$. The object shown in Fig. 1.2a does not satisfy $\sum l_i \vec{u}_i = 0$, and hence is not a polygon. Fig. 1.2b shows a polygon, but its length vector is $L' = (1, 1, 1, 1)$, hence it is not a polygon with length vector L .




Figure 1.2: Non-examples of polygons with length vector $L = (1, 2, 1, 2)$.

In general, the *moduli space* of a system is the space (suitably topologized) of all possible states of that system. This system for us is the set of all polygons with a given length vector L .

When talking about such a moduli space, we do not wish consider certain “redundant” polygons. One redundancy we wish to do away with is that obtained by rotation: if a polygon P

can be obtained from a polygon Q by a rotation of the plane about the origin (recall that all our polygons “start” at the origin), we wish to consider them equal. Since each such equivalence class contains a unique polygon whose last side is oriented in the negative X -direction, getting rid of this redundancy is equivalent to imposing the condition

$$\vec{u}_n = -\vec{e}_1. \quad (1.2)$$

where \vec{e}_1 is the unit length vector in the positive X -direction.

Since we can think of a unit vector in \mathbb{R}^2 as actually lying on the circle S^1 , it follows that a polygon, being a tuple of unit vectors, can be thought of as a point in the n -fold product $S^1 \times \dots \times S^1$. Using Eq. (1.1) and Eq. (1.2), we are now ready to define our moduli space.

Definition 1.3. Let L be a length vector. The **moduli space of polygons with length vector L** , denoted by \mathcal{M}_L , is defined as

$$\mathcal{M}_L = \left\{ (\vec{u}_1, \dots, \vec{u}_n) \in \underbrace{S^1 \times \dots \times S^1}_{n\text{-fold product}} \mid \sum_{i=1}^n l_i \vec{u}_i = 0; \vec{u}_n = -\vec{e}_1 \right\}. \quad (1.3)$$

Thus \mathcal{M}_L is a subspace of the n -torus $T^n = (S^1)^n$. Since rotation in the plane corresponds to the action of the matrix group $SO(2)$, we can also say that

$$\mathcal{M}_L = \left\{ (\vec{u}_1, \dots, \vec{u}_n) \in S^1 \times \dots \times S^1 \mid \sum_{i=1}^n l_i \vec{u}_i = 0 \right\} / SO(2). \quad (1.4)$$

Example 1.3. Let $L = (1, 1, 1)$. Up to rotation, there are only two triangles in the plane with this length vector (starting at the origin, that is). They are shown in Fig. 1.3. Thus \mathcal{M}_L is just two points. In other words, $\mathcal{M}_L = S^0$.

(The reader can verify that no rotation of the plane can bring the triangle in Fig. 1.3a to the one in Fig. 1.3b.)



Figure 1.3: The two triangles in \mathcal{M}_L for $L = (1, 1, 1)$.

Example 1.4. Let $L = (1, 1, 2)$. Then, since $l_1 + l_2 = l_3$, there is precisely one triangle corresponding to L : the degenerate one, obtained by placing V_2 on the edge between V_1 and V_3 . Hence $\mathcal{M}_L = \{\text{a point}\}$.

Example 1.5. Let $L = (1, 1, 3)$. Since $l_1 + l_2 < l_3$, the triangle inequality is not satisfied, so there are no triangles for this length vector. Hence $\mathcal{M}_L = \emptyset$.

As seen in *Example 1.5*, the triangle inequality is a necessary condition for \mathcal{M}_L to be nonempty. This is just a special case of the following:

Proposition 1.1. [*Far08, Lemma 1.1*] *The space \mathcal{M}_L is empty if and only if there exists $i \in [n]$ such that*

$$l_i > \sum_{k \neq i} l_k. \quad (1.5)$$

1.2 Basic results

Let $L = (l_1, \dots, l_n)$ be a length vector. The moduli space \mathcal{M}_L is a subspace of the n -torus, but, as the last examples of the previous section show, it can be of a much lower dimension, and need not be a manifold. So, when is it a manifold? When is it connected? What is its dimension? The results in this section seek to answer these questions.

Definition 1.4. A subset $I \subset [n]$ is called **short with respect to L** if

$$\sum_{i \in I} l_i < \sum_{i \notin I} l_i.$$

Definition 1.5. A subset $I \subset [n]$ is called **long with respect to L** if

$$\sum_{i \in I} l_i > \sum_{i \notin I} l_i.$$

Definition 1.6. A subset $I \subset [n]$ is called **median with respect to L** if

$$\sum_{i \in I} l_i = \sum_{i \notin I} l_i.$$

When L is clear from the context, we drop the “with respect to” and simply say that a subset is short, median or long. Clearly, I is short if and only if its complement $[n] \setminus I$ is long; I is median if and only if $[n] \setminus I$ is also median. Every subset of a short subset is itself short.

Using these new terms, we may restate *Proposition 1.1*: \mathcal{M}_L is empty if and only if there exists a singleton long subset; or, equivalently, \mathcal{M}_L is nonempty if and only if all singletons are short.

While technically the empty set is a short subset, we don’t usually consider it as such (nor do we consider $[n]$ to be long).

Definition 1.7. L is called **generic** if there are no median subsets with respect to L .

Example 1.6. For $L = (1, 1, 1)$, the short subsets are $\{1\}$, $\{2\}$ and $\{3\}$; the long subsets are $\{1, 2\}$, $\{1, 3\}$ and $\{2, 3\}$. There are no median subsets, hence L is generic.

Example 1.7. For $L = (1, 2, 3, 4)$, the short subsets are $\{1\}$, $\{2\}$, $\{3\}$, $\{4\}$, $\{1, 2\}$ and $\{1, 3\}$, with the long subsets being the complements of these. The set $\{2, 3\}$ and (naturally) its complement $\{1, 4\}$ are median subsets. Thus L is not generic.

Remark 1.1. Let $J \subset [n]$ be a median subset that contains n (if not, take its complement). For each $j \in J$, orient \vec{u}_j in the negative X -direction, and for each $k \notin J$, orient \vec{u}_k in the positive X -direction. The result is a degenerate (or collinear) polygon P_J . For example, taking $L = (1, 2, 1, 2)$ and $J = \{3, 4\}$ gives the polygon in Fig. 1.1c. Conversely, any degenerate polygon arises in this way from some median subset J . Thus we conclude that

Median subsets J correspond to degenerate polygons P_J .

Remark 1.2. Suppose $\sigma : [n] \rightarrow [n]$ is a permutation. Then the map

$$\begin{aligned} \phi_\sigma : S^1 \times \dots \times S^1 &\rightarrow S^1 \times \dots \times S^1 \\ (\vec{u}_1, \dots, \vec{u}_n) &\mapsto (\vec{u}_{\sigma(1)}, \dots, \vec{u}_{\sigma(n)}) \end{aligned}$$

is a diffeomorphism from \mathcal{M}_L onto $\mathcal{M}_{L'}$, where $L' = (l_{\sigma(1)}, \dots, l_{\sigma(n)})$. Thus what order the l_i 's appear in does not affect the diffeomorphism type of \mathcal{M}_L . Therefore, when convenient, we will often assume that $l_n \geq l_i$ for all $1 \leq i \leq n$, or that $l_1 \leq \dots \leq l_n$.

We now state, without proof, the main results.

Theorem 1.2. [Far08, Theorem 1.3] *If L is generic, then \mathcal{M}_L is a closed orientable manifold of dimension $n - 3$.*

While we will never have occasion to consider the case when L is not generic, we include the following result here for completeness:

Theorem 1.3. [Far08, Theorem 1.6] *Suppose L is not generic. Then \mathcal{M}_L is a compact $(n - 3)$ -dimensional manifold “with singularities”: it has finitely many singular points which are in one-to-one correspondence with degenerate polygons P_J where $J \subset [n]$ is a median subset containing n .*

Each such P_J has a neighborhood in \mathcal{M}_L homeomorphic to the cone over the product of two spheres

$$S^{|J|-2} \times S^{n-|J|-2}.$$

The next theorem is computationally handy, because it allows us to compute homology groups of \mathcal{M}_L without knowing anything other than L .

Theorem 1.4. [FS07, Theorem 1] *Let $i \in [n]$ be such that $l_i \geq l_j$ for all $j \in [n]$. For each $0 \leq k \leq n - 3$, define*

$$\begin{aligned} a_k &:= \text{no. of short subsets of cardinality } k + 1 \text{ containing } i, \\ b_k &:= \text{no. of median subsets of cardinality } k + 1 \text{ containing } i. \end{aligned}$$

Then, for every $0 \leq k \leq n - 3$, the homology group $H_k(\mathcal{M}_L; \mathbb{Z})$ is free abelian, with

$$\text{Rank } H_k(\mathcal{M}_L; \mathbb{Z}) = a_k + b_k + a_{n-3-k}. \quad (1.6)$$

For the rest of this section, assume that $l_1 \leq \dots \leq l_n$. Unless stated otherwise, we also assume that \mathcal{M}_L is a manifold (that is, L is generic).

We now look at a few results that follow (at least partly) from Theorem 1.4.

Theorem 1.5. \mathcal{M}_L is homeomorphic to the sphere S^{n-3} if and only if the only short subset containing n is the singleton.

Proof. Suppose \mathcal{M}_L is homeomorphic to a sphere. In particular, \mathcal{M}_L is nonempty, so by Proposition 1.1, $\{n\}$ is short. Therefore $a_0 = 1$. We also have

$$\text{Rank } H_k(\mathcal{M}_L; \mathbb{Z}) = \begin{cases} 1, & \text{if } k = 0 \text{ or } n - 3 \\ 0, & \text{otherwise.} \end{cases}$$

Comparing with Eq. (1.6), we must have $a_i = 0$ for $i > 0$. Hence the only short subset containing n is the singleton.

Conversely, suppose the only short subset containing n is the singleton. In particular, the set $\{1, n\}$ is long, so we must have $\vec{u}_1 \neq -\vec{e}_1$. Then the function $f : \mathcal{M}_L \rightarrow \mathbb{R}$ —that maps each polygon to θ_1 , the angle made by the first side with the X -axis—attains a unique maximum and minimum. At any point other than these extrema, $\frac{\partial f}{\partial \theta_1} = 1$, so these are the only critical points. Now we use a theorem of differential topology (see [Mil64, Theorem 1’]): a closed manifold which possesses a smooth real-valued function with exactly two critical points is homeomorphic to a sphere. \square

The proof of the next result has a similar flavor to the previous one, and we omit it here.

Theorem 1.6. \mathcal{M}_L is homeomorphic to the product $S^1 \times S^{n-4}$ if and only if the only short subsets containing n are $\{n\}$ and $\{1, n\}$.

We are now also in a position to answer one of the questions asked earlier: when is \mathcal{M}_L connected?

Lemma 1.7. \mathcal{M}_L is disconnected if and only if $\{n - 2, n - 1\}$ is a long subset. In that case,

$$\text{Rank } H_k(\mathcal{M}_L; \mathbb{Z}) = 2 \binom{n-3}{k}. \quad (1.7)$$

Proof. If $\mathcal{M}_L = \emptyset$, then $\{n\}$ is long, so $\{n - 2, n - 1\}$ is short. So we may assume that \mathcal{M}_L is nonempty.

The rank of the zeroth homology group is $a_0 + a_{n-3}$ (we always assume that \mathcal{M}_L is a manifold, so $b_i = 0$ for all i). Since \mathcal{M}_L is nonempty, $a_0 = 1$. Thus \mathcal{M}_L is disconnected if and only if $a_{n-3} > 0$, that is, there exists at least one short subset of cardinality $n - 2$ containing n . Since we have assumed that $l_1 \leq \dots \leq l_{n-2} \leq l_{n-1} \leq l_n$, that short subset has to be the set $A = \{1, 2, \dots, n - 3, n\}$. Moreover, in that case every subset of A is also short. Thus $a_k = \binom{n-3}{k}$ and $a_{n-3-k} = \binom{n-3}{n-3-k} = \binom{n-3}{k}$, so Eq. (1.7) follows from Eq. (1.6). \square

As one might guess from Eq. (1.7), one has the following theorem:

Theorem 1.8. [KM95, Theorem 1] When the equivalent conditions of Lemma 1.7 hold, \mathcal{M}_L is diffeomorphic to the disjoint union of two $(n - 3)$ -dimensional tori.

We conclude this chapter with some examples.

Example 1.8. Let $L = (1, 1, 1, \dots, 1, n - 1 - \varepsilon)$ for some $0 < \varepsilon \leq 1$. Then $n - 1 - \varepsilon < \sum_{i=1}^{n-1} 1 = n - 1$ and $(n - 1 - \varepsilon) + 1 > \sum_{i=2}^{n-1} 1 = n - 2$, so the only short subset containing n is the singleton. By Theorem 1.5, \mathcal{M}_L is homeomorphic to a sphere.

In particular, if $L = (1, 1, 1, 2)$, \mathcal{M}_L is a circle. This is an example that we will use throughout this thesis to illustrate various ideas.

Example 1.9. Let $L = (1, 2, 2, 2)$. Then $\{2, 3\}$ is a long subset, so by Theorem 1.8, \mathcal{M}_L is a disjoint union of two 1-dimensional tori, that is, $\mathcal{M}_L \simeq S^1 \sqcup S^1$.

Observe that the only short subsets containing 4 are $\{4\}$ and $\{1, 4\}$, so we can even apply Theorem 1.6 here and get $\mathcal{M}_L \simeq S^1 \times S^0$, which is the same as $S^1 \sqcup S^1$.

Example 1.10. Let $L = (1, 1, 1, 1, 1)$. The short subsets containing 5 are $\{5\}$, $\{1, 5\}$, $\{2, 5\}$, $\{3, 5\}$ and $\{4, 5\}$. So $a_0 = 1$, $a_1 = 4$ and $a_2 = 0$. Therefore, by Theorem 1.4, we have

$$\text{Rank } H_0(\mathcal{M}_L; \mathbb{Z}) = a_0 + a_2 = 1,$$

$$\text{Rank } H_1(\mathcal{M}_L; \mathbb{Z}) = a_1 + a_1 = 8,$$

$$\text{Rank } H_2(\mathcal{M}_L; \mathbb{Z}) = a_2 + a_0 = 1.$$

These homology groups suggest that \mathcal{M}_L might be a surface of genus 4, and that is in fact true.

Theorem 1.9. [KM95, Theorem 2] Suppose $n = 5$, \mathcal{M}_L is connected and g is the number of long subsets of cardinality 2. Then \mathcal{M}_L is the closed orientable surface of genus $4 - g$.

Chapter 2

A regular cell structure

When \mathcal{M}_L is a manifold, it admits a regular cell structure that has a succinct combinatorial description purely in terms of L (such a cell structure exists even when \mathcal{M}_L is not a manifold, but we restrict ourselves to the former case here). This cell structure has been described in detail by G. Panina in [Pan12], and the exposition here is based on that paper.

In the current and subsequent chapters, we assume that L is generic and therefore \mathcal{M}_L is a manifold. Unless explicitly mentioned, we also assume that $n \geq 4$.

2.1 Labeling polygons

The first step in this combinatorial construction is to assign a label to each polygon (or, to be precise, equivalence class of polygons). We represent a polygon $P \in \mathcal{M}_L$ by the tuple $(\theta_1, \dots, \theta_n)$, where θ_i is the angle made by the i^{th} side with the horizontal, and $-\pi < \theta_i \leq \pi$ for all i . In other words, since $P = (\vec{u}_1, \dots, \vec{u}_n)$, we have $\vec{u}_i = (\cos \theta_i, \sin \theta_i)$ for all i . The last side is always horizontal and oriented in the negative X -direction, so $\theta_n = \pi$. Before we proceed, a definition:

Definition 2.1. A polygon P is said to be **convex** if the angles $\theta_1, \dots, \theta_n$ are in strictly ascending order, that is, $\theta_1 < \dots < \theta_n$.

The above definition is different from what is usually defined as a convex polygon in \mathbb{R}^2 : all polygons it includes are certainly convex in the usual sense, but it does not include *all* convex polygons, as Fig. 2.1 shows.

Why does our definition of convex imply the usual one? Given a polygon in the plane, each edge of the polygon divides the plane into two half-planes. A sufficient condition for the polygon to be convex is that, for each of its edges, it lie entirely within one of these half-planes. If the angles made by the edges with the horizontal are in increasing order, then the polygon must lie entirely within one half-plane of each edge, and hence is convex.

We can now continue. For assigning labels, it will be convenient to split the polygons into two classes—those with and those without parallel edges—and to deal with these classes individually:




Figure 2.1: These pictures show when our definition of convex coincides with the usual one, and when it doesn't.

1. Polygons without parallel edges Given a polygon P , no two of whose edges are parallel (that is, $\theta_i \neq \theta_j$ if $i \neq j$; note that we *are* allowing $\theta_i = -\theta_j$, so the “parallel” here is in a very strict sense), there is a way to associate a unique convex polygon to P : arrange the angles θ_i in ascending order. This gives rise to a permutation $\lambda \in S_n$, since no two edges are parallel. That is, λ is the unique permutation such that $\theta_{\lambda(1)} \leq \dots \leq \theta_{\lambda(n)}$. By hypothesis, $\lambda(n) = n$. Now construct the polygon P^{conv} whose i^{th} edge has length $l_{\lambda(i)}$ and makes an angle of $u_{\lambda(i)}$ with the horizontal. This construction is made clear with an example in Fig. 2.2.

Remark 2.1. The polygon P^{conv} itself is not an element of \mathcal{M}_L . In fact, $P^{conv} \in \mathcal{M}_{\lambda L}$.

The label for P is just $\lambda([n])$, but we want to express it in a way that will allow us to deal with the situation when our polygons *do* have parallel edges. We therefore make the following definitions:

Definition 2.2. An **ordered partition**

$$\alpha = (I_1 \dots I_k)$$

of $[n]$ is a partition in which the order of the I_j 's matters. (Here $I_j \subset [n]$; $I_i \cap I_j = \emptyset$ for $i \neq j$; and $\cup_j I_j = [n]$.) The subsets I_j 's are called **blocks** of α .

Note that, if


$$\beta = (I_2 I_1 I_3 I_4 \dots I_k)$$

is another partition of $[n]$ with the same sets but in a different order, then $\beta \neq \alpha$.


Definition 2.3. A **cyclically ordered partition** of $[n]$ is an ordered partition which is equivalent to any partition obtained from it by a cyclic permutation of its blocks. That is, if $\alpha = (I_1 \dots I_k)$ is a cyclically ordered partition, then $\alpha = (I_2 \dots I_k I_1) = (I_3 \dots I_k I_1 I_2) = \dots = (I_k I_1 \dots I_{k-1})$. When dealing with such partitions, we will always assume that the set containing n is the last set.

Definition 2.4. Fix a length vector L . A cyclically ordered partition of $[n]$ is called **admissible** if each block in the partition is a short subset of $[n]$.


Example 2.1. For $L = (1, 1, 1, 2)$, the ordered partition $(\{1\}\{2, 3\}\{4\})$ is different from $(\{1, 2\}\{3\}\{4\})$, because they have different blocks. It is also different from $(\{2, 3\}\{1\}\{4\})$, because their blocks



(a) The polygon P . Observe that no two edges point in the same direction.



(b) The edges of P drawn originating from the same point. Arranged in ascending order, they are 2,4,1,3,5. Thus
 $\lambda(1) = 2$
 $\lambda(2) = 4$
 $\lambda(3) = 1$
 $\lambda(4) = 3$
 $\lambda(5) = 5$



(c) The polygon P^{conv} . Its first edge is $\lambda(1) = 2$, that is, the edge that was the second edge in P , and similarly for the rest of the edges. Clearly, P^{conv} is convex. The label for P is $(\{2\}\{4\}\{1\}\{3\}\{5\})$.

Figure 2.2: Labeling a polygon without parallel edges

are in different orders.

The cyclically ordered partition $(\{1,2\}\{3\}\{4\})$ is, however, equal to the cyclically ordered partition $\{4\}(\{1,2\}\{3\})$ and, as mentioned above, we will always write it with $\{4\}$ as the last block.

Further, the ordered partition $(\{1\}\{2\}\{3,4\})$ is not admissible, because $\{3,4\}$ is not a short

subset, while $(\{1, 2\}\{3\}\{4\})$ is admissible.

The **label** assigned to a polygon P is the admissible partition consisting of singletons whose order is given by λ . In other words, the label is $(\{\lambda(1)\}\{\lambda(2)\} \dots \{\lambda(n-1)\}\{n\})$.

Fig. 2.2c shows how this label is assigned. We point out here that, since for a polygon without parallel edges the label consists only of singletons, it follows that we don't have to check whether the partition is admissible (because all singletons are short subsets). In other words, any partition of $[n]$ consisting of singletons is a label. Assigning a cyclically ordered partition is justified because applying a cyclic permutation to the sets in the partition won't change the ordering of the edges. The fact that the n -set is written last is just a convention we have adopted.

There is an obvious way to recover P from P^{conv} : the i^{th} edge of P is the $\lambda^{-1}(i)^{th}$ edge of P^{conv} , so the first edge (in the figure) is the $\lambda^{-1}(1) = 3^{rd}$ edge and so on. Thus, given a label λ , we have established a bijective correspondence

$$\begin{aligned} \{\text{Polygons in } \mathcal{M}_L \text{ labeled by } \lambda\} &\iff \{\text{Convex polygons in } \mathcal{M}_{\lambda L}\} \\ P &\leftrightarrow P^{conv}. \end{aligned} \quad (2.1)$$

In fact, since perturbing the θ_i 's of P slightly ends up in a tiny perturbation of the resulting P^{conv} , and conversely, we end up with

Lemma 2.1. *Given an ordered partition λ of $[n]$ into singletons, the subspace of \mathcal{M}_L consisting of all polygons labeled by λ is homeomorphic to the space of convex polygons in $\mathcal{M}_{\lambda L}$.*

We now turn to the other class:

2. Polygons with parallel edges For polygons with parallel edges, we proceed the same way as we did for the earlier polygons: arrange the edges in order of increasing angle and then look at the resulting convex polygon. However, the problem here is that for each pair of angles $\theta_i = \theta_j$, it is not clear which one should be placed first in the resulting ascending order. The solution is to "glue" all such edges together, and consider them as constituting one edge.

For example, consider a polygon P that has 8 edges, of which edges 1, 3 and 7 are parallel, and edges 2 and 4 are parallel, with the rest of the edges all being in different directions. Then we think of the convex polygon P^{conv} obtained from this polygon as one with 5 edges, with one of the edges consisting of the original edges 1, 3 and 7 glued together, and another consisting of the original 2 and 4 glued together. Fig. 2.3 makes this clear.

The label for such a polygon again consists of an admissible partition, but now the sets are not all singletons. Parallel edges all end up in the same set in such a partition. The construction of the partition, however, is more or less the same as in the first case: starting with the least angle, proceed counterclockwise and place edges with different angles in different sets of the partition (see Fig. 2.3b). So, for P , we have the partition $(\{5\}\{2, 4\}\{6\}\{1, 3, 7\}\{8\})$. This is the label for P .

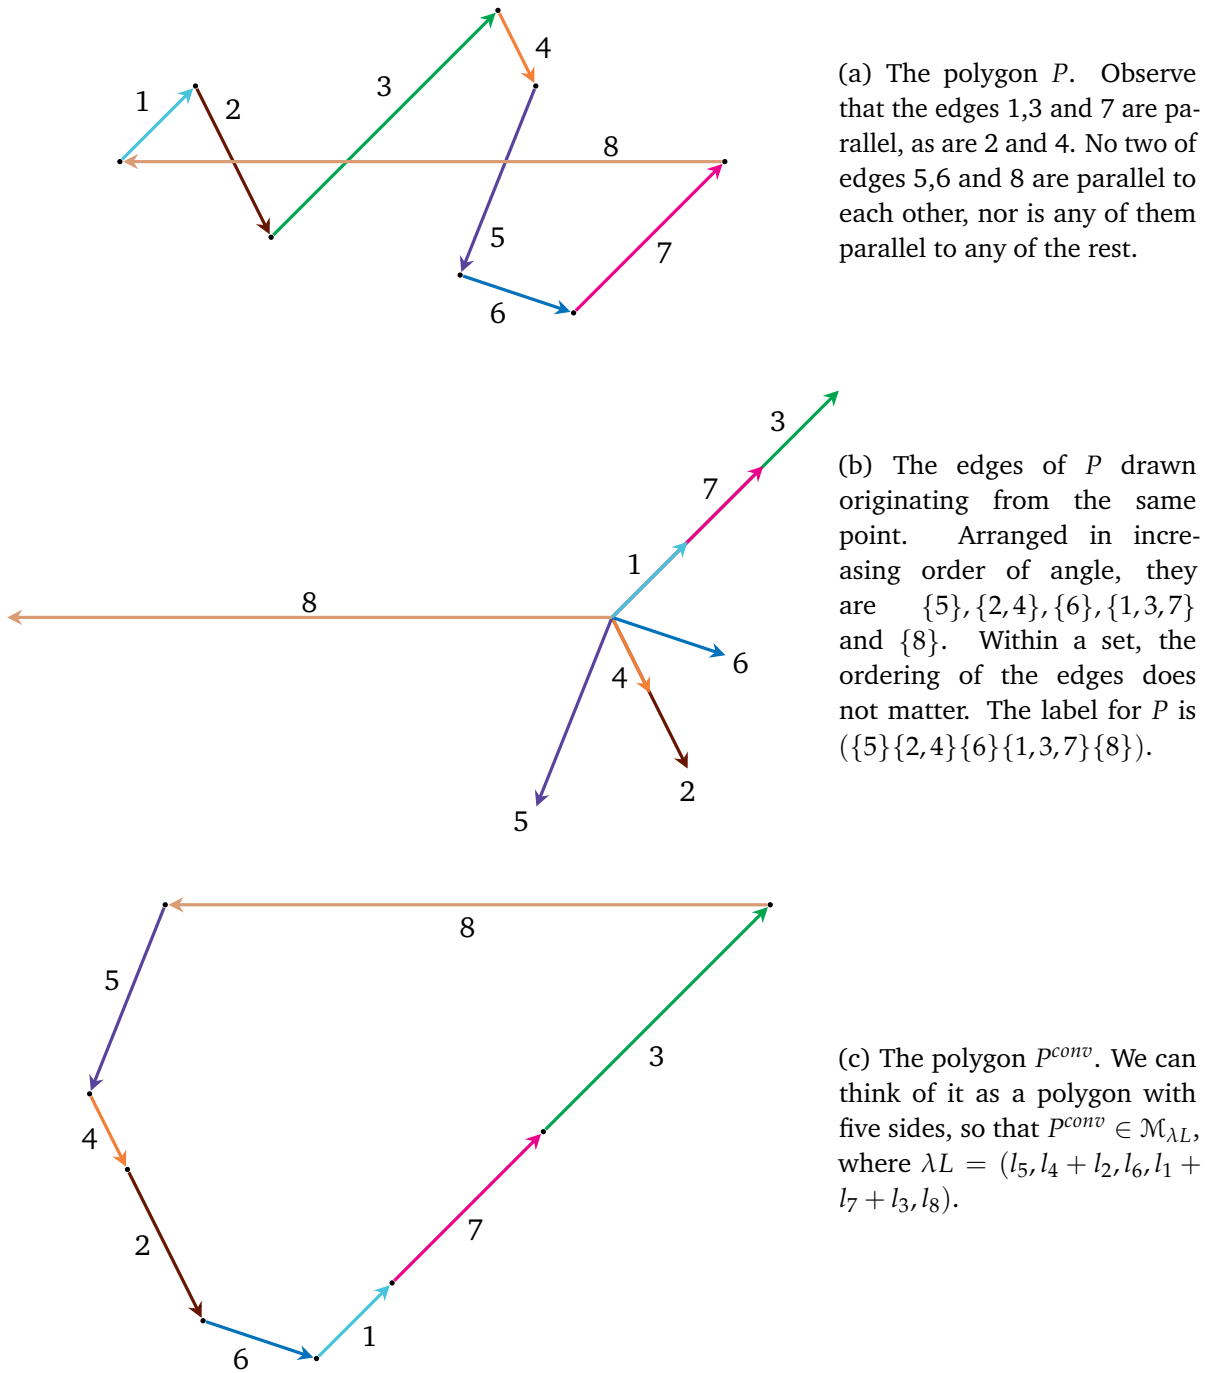


Figure 2.3: Labeling a polygon with parallel edges

In general, if we end up with an admissible partition $\lambda = (I_1 I_2 \dots I_k)$ consisting of k subsets of $[n]$, then we can think of the resulting polygon as having k edges and residing in $\mathcal{M}_{\lambda L}$, where

$$\lambda L = (\tilde{l}_1, \tilde{l}_2, \dots, \tilde{l}_k), \text{ with } \tilde{l}_j = \sum_{i \in I_j} l_i.$$

Fig. 2.3c shows how this is done for P .

As before, we can recover our original polygon using λ and the resulting convex polygon. For the i^{th} edge of the polygon, suppose the set I_j in λ contains i . Then this edge points in the direction that the j^{th} edge of the convex polygon points, and has length l_j . For example, for P , 4 lies in the set I_2 of λ . So the fourth edge of P has length l_4 and points in the direction that the second edge of P^{conv} points. So, similar to the polygons with no parallel edges, we have a bijection and, in fact, a homeomorphism:

Lemma 2.2. *Given an admissible partition λ of $[n]$ into k blocks, the subspace of \mathcal{M}_L consisting of all polygons labeled by λ is homeomorphic to the space of convex k -gons in $\mathcal{M}_{\lambda L}$.*

We recover Lemma 2.1 when $k = n$.

2.2 The cells

In this section, we come to the meat of this chapter: describing the cells of the complex. Essentially, a cell is formed by polygons having the same label. The next few results make this precise.

Denote by $\mathcal{M}_L^{\text{conv}}$ the set of all convex (in the sense of Definition 2.1) polygons in \mathcal{M}_L .

Lemma 2.3. [Pan12, Lemma 1.2]

- (1) $\mathcal{M}_L^{\text{conv}}$ is an open subset of \mathcal{M}_L homeomorphic to the open $(n - 3)$ -dimensional ball.
- (2) The closure $\overline{\mathcal{M}_L^{\text{conv}}}$ is homeomorphic to the closed $(n - 3)$ -dimensional ball.

Proof. (1) The proof proceeds by induction on n , the number of components of L . We first analyze the case $n = 3$, because it is instructive.

For $n = 3$, there are two possible polygons (L is generic), of which one is non-convex (Fig. 2.4a) and one is convex (Fig. 2.4b). So $\mathcal{M}_L^{\text{conv}}$ is just a point.



Figure 2.4: The two possible triangles for $n = 3$

For $n = 4$, consider the map that sends a polygon to θ_1 , the angle made by the first edge with the horizontal. The image of this function is either a proper subset of the interval $(-\pi, \pi]$, or is the whole circle. In the former case, since \mathcal{M}_L is compact, this function has a maximum and a minimum. Since L is generic, the image is not just a point, so these are distinct. Let a_1 be the minimum value of this function. In the latter case (when the image is a circle), let a_1 be $-\pi$. For a particular polygon to be convex, none of the interior angles can be greater than π , so

in particular $\theta_i < \theta_{i+1} < \theta_i + \pi$ for all $1 \leq i < n$ and also $0 = \theta_n - \pi < \theta_1 + \pi < \theta_n = \pi$. So on \mathcal{M}_L^{conv} , we must have $a_1 < \theta_1 < 0$. However, this is a necessary condition, not a sufficient one. We must also ensure that the triangle formed by the last three vertices is convex, as shown in Fig. 2.5.



Figure 2.5: That $\theta_1 < 0$ is a necessary condition, not sufficient: the triangle formed by the second and third edge with the dashed line must also be convex.

So far, we have ascertained that a convex 4-gon must have its first angle between a_1 and 0, and must be in such a position so as to ensure that the triangle formed by the last three vertices is also convex. However, even these conditions are not sufficient. It is possible that the triangle is convex but the second angle is less than θ_1 (Fig. 2.6a).




Figure 2.6: The triangle formed by the last three vertices may be convex, but the 4-gon itself is convex only when $\theta_1 < b_1$.

So we need to ask: what is the greatest value of θ_1 for which one is guaranteed to have a convex 4-gon? Consider Fig. 2.6b, where the first and second sides are parallel. Call the value of θ_1 for this polygon b_1 . At any θ_1 greater than b_1 , we cannot have a convex 4-gon (Fig. 2.6a). At any θ_1 strictly between a_1 and b_1 , we have exactly one convex 4-gon (Fig. 2.6c). We therefore have a bijection

$$\begin{aligned} \mathcal{M}_L^{conv} &\rightarrow (a_1, b_1) \\ P &\mapsto \theta_1 \end{aligned}$$

which is a homeomorphism. The proof for $n = 4$ is thus complete.

Now suppose that, for some $n \geq 4$, the space \mathcal{M}_L^{conv} is homeomorphic to I^{n-3} , the product of $n - 3$ open intervals. We need to show that, for $n + 1$, the space \mathcal{M}_L^{conv} is homeomorphic to I^{n-2} , the product of $n - 2$ open intervals. We proceed exactly as we did for $n = 4$. Let a_1 be the minimum possible value of θ_1 , and let $b_1 < 0$ be the value of θ_1 for which $\vec{u}_2 = \vec{u}_3 = \dots = \vec{u}_{n-2}$. Then, for each value of θ_1 between a_1 and b_1 , the space \mathcal{M}_L^{conv} for the last n vertices is I^{n-3} , and conversely, for each convex $(n + 1)$ -gon, the first angle must lie between a_1 and b_1 , and the n -gon of the last n vertices, P' , must also be convex. In short, we have a homeomorphism

$$\begin{aligned} \mathcal{M}_L^{conv} &\rightarrow (a_1, b_1) \times I^{n-3} \simeq I^{n-2} \\ P &\mapsto (\theta_1, P'), \end{aligned}$$

and the proof is complete, because the product of k open intervals is homeomorphic to the open k -ball (for $k > 0$).

(2) The proof for the second statement goes precisely the same as the first, except that each instance of the open interval (a_1, b_1) is replaced by the closed interval $[a_1, b_1]$. (For example, the 4-gon in Fig. 2.6b lies in the closure of \mathcal{M}_L^{conv} , because it can be approximated by polygons of the type shown in Fig. 2.6c.) \square

Lemma 2.4. [Pan12, Lemma 2.3] *Given an ordered partition λ of the set $[n]$ into n nonempty blocks, the subset of \mathcal{M}_L of all polygons labeled by λ is an open $(n - 3)$ -ball.*

Proof. By Lemma 2.1, the space of polygons labeled by λ is homeomorphic to the space $\mathcal{M}_{\lambda L}^{conv}$, which by Lemma 2.3 (1) is homeomorphic to an open $(n - 3)$ -ball. \square

Lemma 2.5. [Pan12, Lemma 2.4] *Given an admissible partition λ of the set $[n]$ into k nonempty blocks, the subset of \mathcal{M}_L of all polygons labeled by λ is an open $(k - 3)$ -ball.*

Proof. By Lemma 2.2, the space of polygons labeled by λ is homeomorphic to the space $\mathcal{M}_{\lambda L}^{conv}$ of convex k -gons in $\mathcal{M}_{\lambda L}$, which by Lemma 2.3 (1) is homeomorphic to an open $(k - 3)$ -ball. \square

The upshot of the above results is that each admissible partition of $[n]$ into k blocks gives us a label, and that each such label corresponds to an open k -ball in \mathcal{M}_L . Since each polygon is assigned a unique label, the union of these balls is the whole \mathcal{M}_L (note that not all of these “open” balls are open in \mathcal{M}_L , only the ones of dimension $n - 3$).

Definition 2.5. A **regular CW complex** is a CW complex in which the attaching maps are all injective.



Figure 2.7: Examples of regular and non-regular cell complexes.

Example 2.2. The cell complex obtained by attaching a 1-cell to a 0-cell is not regular, because the attaching map takes the boundary of the 1-cell (which is S^0) to a single point, hence it is not an injection (Fig. 2.7a). The cell complex consisting of two 1-cells and two 0-cells, with each 1-cell having both 0-cells in its boundary, is regular (Fig. 2.7b).

The next theorem, the main result of this chapter, states that the $(k - 3)$ -balls of Lemma 2.5 are in fact the cells of a regular CW complex on \mathcal{M}_L . To state it precisely, we clarify some terminology. An *open cell* is the set of all polygons having the same label. It is an open ball, and its closure is a closed ball, which we will call a *closed cell*. For a cell C , either closed or open, its label $\lambda(C)$ is defined to be the label of its interior points.

Definition 2.6. A label λ is said to **refine** or be **finer than** a label λ' if λ is a refinement of λ' as a partition (that is, each set of λ is contained in a set of λ') and if the ordering of λ is inherited from that of λ' . Since our partitions are cyclically ordered and we always require the set containing n to be last in an ordered partition, a refinement of such a partition may contain a subset of the n -set as the first set. The next example illustrates this.

Example 2.3. Let $\lambda = (\{1,2\}\{4\}\{3,5\})$. The label $(\{1\}\{2\}\{4\}\{3,5\})$ refines λ , as does $(\{2\}\{1\}\{4\}\{3,5\})$, while the label $(\{4\}\{1\}\{2\}\{3,5\})$ does not refine λ , because the set $\{4\}$ must appear *after* all subsets of $\{1,2\}$.

The label $(\{1,2\}\{4\}\{3\}\{5\})$ refines λ , as does $(\{1,2\}\{4\}\{5\}\{3\})$ but, according to our convention, the latter is written as $(\{3\}\{1,2\}\{4\}\{5\})$. So it is possible for a label to refine another while having a subset of the n -set as its first set.

Theorem 2.6. [Pan12, Theorem 2.6] *The moduli space \mathcal{M}_L admits a regular CW complex structure \mathcal{K}_L . Its complete combinatorial description is as follows:*

- (1) *The k -cells are labeled by admissible partitions of the set $[n]$ into $(k + 3)$ nonempty parts.*
- (2) *A closed cell C belongs to the boundary of some closed cell C' if and only if the partition $\lambda(C')$ is finer than $\lambda(C)$.*

Proof. As we've already seen, polygons sharing a label (a partition of $[n]$ into k parts) form an open $(k - 3)$ -ball, so (1) follows. One can infer from the proof of Lemma 2.3 that a polygon P belongs to the boundary of a closed cell C if and only if it is obtained by some subset of the edges (which are not parallel in C) becoming parallel in P . This is precisely the assertion that the label P is refined by $\lambda(C)$, and (2) follows.

The attaching maps of the CW complex follow immediately: a k -gon in the boundary of a cell C maps to the (necessarily unique) j -gon (where $j < k$) obtained by gluing together the edges which have become parallel in the k -gon. Conversely, given such a j -gon in the $(k - 1)$ -skeleton of the complex, there is a unique k -gon that maps to it, namely the one obtained by inserting vertices in the j -gon to ensure that the resulting polygon has the label $\lambda(C)$. The regularity of the complex is now evident. \square

In general, we will often identify a cell in the complex \mathcal{K}_L with its label, and use both interchangeably.

Example 2.4. Let $L = (1, 1, 1, 2)$. As we know (Example 1.8), \mathcal{M}_L in this case is a circle. The 0-cells for \mathcal{K}_L are given by admissible partitions of $[4]$ into 3 sets. These are:

1. $(\{1,2\}\{3\}\{4\})$
2. $(\{1\}\{2,3\}\{4\})$
3. $(\{1,3\}\{2\}\{4\})$
4. $(\{3\}\{1,2\}\{4\})$
5. $(\{2,3\}\{1\}\{4\})$




Figure 2.8: \mathcal{K}_L for when M_L is a circle.

6. $(\{2\}\{1,3\}\{4\})$

The 1-cells are given by admissible partitions of $[4]$ into singletons. These are:

1. $(\{1\}\{2\}\{3\}\{4\})$
2. $(\{1\}\{3\}\{2\}\{4\})$
3. $(\{2\}\{1\}\{3\}\{4\})$
4. $(\{2\}\{3\}\{1\}\{4\})$
5. $(\{3\}\{1\}\{2\}\{4\})$
6. $(\{3\}\{2\}\{1\}\{4\})$

A 1-cell contains in its boundary all 0-cells it refines. For example, the boundary of the 1-cell $(\{1\}\{2\}\{3\}\{4\})$ consists of the 0-cells $(\{1,2\}\{3\}\{4\})$ and $(\{1\}\{2,3\}\{4\})$. Fig. 2.8 illustrates the cellular decomposition of \mathcal{M}_L .

Example 2.5. Let $L = (1, 2, 2, 2)$. Here \mathcal{M}_L is a disjoint union of circles (Example 1.9), and the short subsets are the singletons and all sets of cardinality two containing 1. We do not list all admissible partitions here, but just refer to Fig. 2.9, which shows the cellular decomposition.

Example 2.6. Let $L = (1, 1, 1, 1, 1)$. In this case M_L is a surface of genus 4 (Theorem 1.9). The top-dimensional cells are labeled by admissible partitions of $[5]$ into singletons. There are exactly $4! = 24$ of these. Since any two-element subset is short, each top-dimensional cell refines five admissible partitions, and hence is a pentagon. We conclude that \mathcal{K}_L consists of 24 pentagons patched together.




Figure 2.9: \mathcal{K}_L for when M_L is a disjoint union of circles. We have omitted the braces in the labels, writing elements of a set contiguously.

2.3 Permutohedra

As promised, we have delivered a CW complex on \mathcal{M}_L that can be combinatorially described. But we can take this further. This section introduces a type of object called a *permutohedron* that will be used in the next section to complete our task.

Definition 2.7. The **permutohedron** Π_n ($n \geq 2$) is defined as the convex hull of all points in \mathbb{R}^n that are obtained by permuting the coordinates of the point $(1, 2, \dots, n)$.

Example 2.7. The permutohedron Π_2 is the convex hull of $(1, 2)$ and $(2, 1)$ in \mathbb{R}^2 . It is just the line segment joining the two points.

Example 2.8. The permutohedron Π_3 is the convex hull of the points $(1, 2, 3)$, $(1, 3, 2)$, $(2, 1, 3)$, $(2, 3, 1)$, $(3, 1, 2)$ and $(3, 2, 1)$ in \mathbb{R}^3 . In other words, it is a hexagon lying on the plane $x + y + z = 6$ with the aforementioned points as vertices. The vertices of this hexagon can be identified with elements of S_3 , the symmetric group on 3 letters. The edges can be identified with ordered partitions of the set $\{1, 2, 3\}$ into two sets. These are depicted in Fig. 2.10.

The similarity between Fig. 2.8 and Fig. 2.10 is not coincidental, as we will soon see.

Example 2.9. The permutohedron Π_4 is a three-dimensional polyhedron with 24 vertices. We do not depict it here, but pictures abound on the internet. The reader may visualize it as a truncated octahedron: take an octahedron and remove a square pyramid from around each vertex (of which there are six) such that no two such square pyramids intersect.

The vertices of Π_4 correspond to the elements of S_4 , the edges (of which there are 36) to partitions of $\{1, 2, 3, 4\}$ into three parts, and the (8 hexagonal and 6 square) faces to partitions into two parts.

Definition 2.8. A **polytope** is the convex hull of finitely many points in \mathbb{R}^d .

Following [Pan12], we summarize here some facts about permutohedra. We provide no proofs but trust that the above examples illuminate matters considerably:

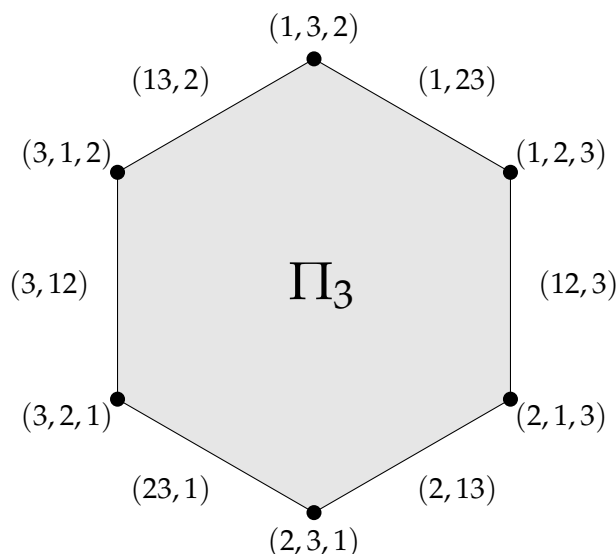


Figure 2.10: The permutohedron Π_3 is a 2-dimensional polytope whose boundary is a hexagon.

1. Π_n is an $(n - 1)$ -dimensional polytope.
2. The k -dimensional faces of Π_n are labeled by ordered partitions of $[n]$ into $(n - k)$ nonempty parts. In particular, the vertices are labeled by elements of S_n . (The term *labeled by* is used rather suggestively, because we are immediately going to associate these with the labels we have already defined for cells of \mathcal{M}_L)
3. A face F' of Π_n is contained in (the boundary of) a face F if and only if the label of F' is finer than the label of F . Refinement here is as described in Definition 2.6.
4. A face of Π_n is the Cartesian product of permutohedra of smaller dimensions. (For example, each face of Π_3 is a line segment (Π_2); each face of Π_4 is a line segment (Π_2), a hexagon (Π_3) or a square ($\Pi_2 \times \Pi_2$).)

Before we can connect permutohedra to the ideas developed in the previous section, we need to talk about face posets.

2.4 Face posets and dual complexes

This section describes how the faces of the complex \mathcal{K}_L can be described in terms of permutohedra. This description naturally leads us to consider the dual cell structure \mathcal{K}_L^* , and we discuss that briefly. But first, some definitions.

Definition 2.9. The **face poset** of a regular CW complex is the poset (P, \prec) obtained by ordering its cells by inclusion: $\alpha \prec \beta$ if and only if α is contained in the boundary of β .

Since we can draw a Hasse diagram of any face poset, the Hasse diagram for the face poset of the complex \mathcal{K}_L of Fig. 2.8 is depicted in Fig. 2.11.




Figure 2.11: The Hasse diagram of the face poset for the cell complex \mathcal{K}_L for $L = (1, 1, 1, 2)$. Since the set containing 4 is always a singleton, we have omitted it from all labels.

We will need some ideas and facts from the theory of polytopes, and we source most of these from G. Ziegler’s ([Zie95]).

Definition 2.10. Two polytopes are said to be **combinatorially equivalent** if their face posets are isomorphic; that is, there is an order-preserving bijection between their face sets.

Definition 2.11. Two polytopes are said to be **combinatorially dual** if their face posets are anti-isomorphic; that is, there is an order-reversing bijection between their face sets. A polytope Q that is combinatorially dual to the polytope P is said to be the **dual of P** and is denoted by P^* (we say “the” dual, but it is determined only up to combinatorial equivalence).

It can be shown that the dual of a polytope always exists (for example, using the construction of a polar polytope outlined in [Zie95]).

Polytopes have a unique face of maximum dimension (the polytope itself), and a unique (-1) -dimensional face (that is, the empty face), so that the face poset of a polytope becomes a lattice, called the *face lattice*. The maximum and minimum dimensional faces are denoted by $\hat{1}$ and $\hat{0}$ respectively. For each face F of such a lattice, the interval $[\hat{0}, F]$ is isomorphic to the face lattice of F .

Definition 2.12. For each face F of the polytope P , the interval $[F, \hat{1}]$ is isomorphic to the face lattice of a polytope, called the **face figure** of F , and denoted by P/F . It is called a **vertex** (respectively, **edge**) **figure** if F is a vertex (respectively, edge).

If the polytope P is dual to the polytope Q with the face F mapping to G under the anti-isomorphism, then G is the face figure P/F of F . In general,

$$\dim(P/F) = \dim(P) - \dim(F) - 1.$$

Example 2.10. For a regular polyhedron, the vertex figure may be constructed by cutting off a small slice containing the vertex (that is, taking a hyperplane that passes through the polytope with the given vertex on one side of the hyperplane and all other vertices on the other side) and looking at the exposed hypersurface. Take an octahedron, for example. Cut off a small square pyramid containing a vertex. The exposed surface is a square. This is the vertex figure of the vertex we just cut off. As described above, the vertex figure can also be obtained by going to the dual polyhedron (the cube) and looking at the corresponding codimension-1 face. Again, this is a square.

The notion of a face figure can be extended to posets of CW complexes, but now, since we do not necessarily have lattices, the resulting face figures could be objects that do not have a CW decomposition at all.

Definition 2.13. For a cell C of a CW complex \mathcal{K} , an object (not necessarily a CW complex) having the same face poset as the interval $[C, \hat{1})$ (or $[C, \hat{1}]$ if $\hat{1}$ exists) is called the **face figure** of C . Analogous definitions hold for **vertex figure** and **edge figure**.

Using the vocabulary developed in the previous and current sections, we are ready to state our result. Recall that \mathcal{K}_L is the regular cell decomposition of \mathcal{M}_L .

Proposition 2.7. [Pan12, Proposition 2.7]

- (1) The vertex figure of any vertex v of \mathcal{K}_L is combinatorially dual to the Cartesian product of three permutohedra.
- (2) The face figure of any k -dimensional cell is combinatorially dual to the Cartesian product of $(k + 3)$ permutohedra.

Proof. (1) By Theorem 2.6, we know that the label of a 0-dimensional cell F is a partition of $[n]$ into three sets, say A, B and C , with cardinalities a, b and c respectively. Consider the interval $[F, \hat{1})$. Any cell in this interval has a label which consists of a partition of A followed by a partition of B followed by a partition of C . Thus, it follows that $[F, \hat{1})$ is isomorphic to the product of the posets of cyclically ordered partitions of A, B and C . The poset of cyclically ordered partitions of A , however, is combinatorially dual to the face poset of Π_a , and similarly for B and C , so the result follows.

(2) is just an extension of (1) to the case when the label has more than 3 parts. □

Example 2.11. For the \mathcal{K}_L of Fig. 2.8, the $[F, \hat{1})$ for an arbitrary 0-cell is shown in Fig. 2.12, as are the actual vertex figure and the product of permutohedra it is dual to.




Figure 2.12: The vertex figure of an arbitrary vertex of \mathcal{K}_L for $M_L = S^1$ is combinatorially dual to the product $\Pi_2 \times \Pi_1 \times \Pi_1$ of three permutohedra.

With this result, we have described completely the combinatorial structure of M_L .

Chapter 3

Discrete Morse theory

3.1 A primer on discrete Morse theory

This section serves as a brief introduction to discrete Morse theory. It closely follows R. Forman's treatment in [For02].

Discrete Morse theory is a technique for analyzing the topology of a CW complex by defining a special type of function on it, called a *discrete Morse function*. The reasoning for this terminology is sound: the “discrete” is because we do not assign a continuous set of values to each point in the space, but only a discrete set of values to the cells in the complex; the “Morse theory” part is justified because, similar to classical Morse theory, we have notions of critical points and gradient paths which can be used to state discrete versions of the Morse inequalities. More will become clear as definitions and examples are presented. We will, unless mentioned otherwise, always assume that all spaces are manifolds and all complexes are regular CW complexes.

Let \mathcal{K} be a finite regular CW complex. For a cell α of \mathcal{K} , $\alpha^{(p)}$ indicates that it is a p -cell, and both $\alpha^{(p)} < \beta^{(p+1)}$ and $\beta^{(p+1)} > \alpha^{(p)}$ indicate that α lies in the boundary of β .

Definition 3.1. A function

$$f : \mathcal{K} \rightarrow \mathbb{R}$$

is called a **discrete Morse function** if, for every $\alpha^{(p)} \in \mathcal{K}$, the following hold:

1. $s(\alpha) := \left| \left\{ \beta^{(p+1)} > \alpha \mid f(\beta) \leq f(\alpha) \right\} \right| \leq 1$, and
2. $i(\alpha) := \left| \left\{ \gamma^{(p-1)} < \alpha \mid f(\gamma) \geq f(\alpha) \right\} \right| \leq 1$.

So, a discrete Morse function, in the words of Forman, “roughly speaking, assigns higher numbers to higher dimensional [cells], with at most one exception, locally, at each [cell].” The following lemma, which is a direct consequence of the definition, makes precise the “at most one exception locally”, and is useful when one wants to show that a particular function is not a discrete Morse function:

Lemma 3.1. [For98, Lemma 2.5] Suppose $p > 0$. For each cell $\alpha^{(p)}$ of \mathcal{K} , at least one of the inequalities in Definition 3.1 must be strict. In other words, $s(\alpha) + i(\alpha) \leq 1$.

Proof. Suppose there is some $\alpha^{(p)}$ for which neither inequality is strict. That is, there exists a $\beta^{(p+1)} > \alpha$ and a $\gamma^{(p-1)} < \alpha$ such that $f(\beta) \leq f(\alpha)$ and $f(\gamma) \geq f(\alpha)$. But that means

$$f(\gamma) \geq f(\beta). \tag{3.1}$$

Now, since we are in a regular complex, there exists some $\alpha'^{(p)}$ such that $\gamma < \alpha' < \beta$. Applying condition 1 of the definition to γ , we see that $f(\alpha') > f(\gamma)$. Applying condition 2 to β , we see that $f(\alpha') < f(\beta)$. But these inequalities contradict Eq. (3.1). \square

When both inequalities are strict (that is, when $s(\alpha) = i(\alpha) = 0$), we get an analog of the critical points in classical Morse theory:

Definition 3.2. A cell $\alpha^{(p)}$ is called a **critical cell** or a **critical p -cell** if the following hold:

1. $f(\beta) > f(\alpha)$ for all $\beta^{(p+1)} > \alpha$, and
2. $f(\gamma) < f(\alpha)$ for all $\gamma^{(p-1)} < \alpha$.

We now demonstrate some of the definitions introduced above with an example:

Example 3.1. Consider the regular complex \mathcal{K}_L when $L = (1, 1, 1, 2)$. Recall that here \mathcal{M}_L is a circle (Example 1.8). Fig. 3.1 shows \mathcal{K}_L .




Figure 3.1: \mathcal{K}_L for when \mathcal{M}_L is a circle. We omit the braces and 4-set in the cell labels.

Now, since a discrete Morse function assigns a number to each cell, we again show the cell decomposition, but now we replace the cell labels by the value assigned to each cell.

The assignment of values shown in Fig. 3.2 would constitute a discrete Morse function but for the fact that $i((3,2,1)) = 2$. To rectify this, we can just increase the value assigned to it. Fig. 3.3 shows this modified function f , which is indeed a discrete Morse function.




Figure 3.2: A non-example of a discrete Morse function




Figure 3.3: The discrete Morse function f

What are the critical cells of f ? First, we look at all 1-cells β with $i(\beta) = 0$. These are $(3, 1, 2)$, $(2, 3, 1)$ and $(1, 2, 3)$. Next, we look at all 0-cells α with $s(\alpha) = 0$. These are $(3, 12)$, $(2, 13)$ and $(1, 23)$. Thus we end up with three critical 0-cells and three critical 1-cells.

It is no coincidence that a circle is homotopy equivalent to a CW complex with three 0-cells and three 1-cells.

Theorem 3.2. [For98, Corollary 3.5] Suppose \mathcal{K} has a discrete Morse function. Then \mathcal{K} is homotopy equivalent to a CW complex with exactly one cell of dimension p for each critical p -cell.


A proof of this theorem may be looked up in R. Forman's [For98], and a discussion about why

the proof works is present in [For02]. The essential idea is that of an *elementary collapse*, which we now illustrate.

Definition 3.3. For a discrete Morse function $f: \mathcal{K} \rightarrow \mathbb{R}$, the **level sub-complex** $\mathcal{K}(c)$ is defined to be the sub-complex containing all cells α such that $f(\alpha) \leq c$, and all the cells in their boundaries. That is

$$\mathcal{K}(c) = \bigcup_{f(\alpha) \leq c} \cup_{\gamma \leq \alpha} \gamma.$$

Some level sub-complexes of the discrete Morse function f in Fig. 3.3 are shown in Fig. 3.4.



(a) The level sub-complex $\mathcal{K}(1)$ (b) The level sub-complex $\mathcal{K}(3)$ (c) The level sub-complex $\mathcal{K}(4)$

Figure 3.4: Some level sub-complexes for f on \mathcal{K} .

Observe that, when we move from $\mathcal{K}(1)$ to $\mathcal{K}(3)$, we add two non-critical cells, and not just any two cells: a pair $\alpha^{(p)}$ and $\beta^{(p+1)}$ such that $\alpha < \beta$ and $f(\alpha) \geq f(\beta)$. This is in general what happens when $f^{-1}(a, b]$ has no critical cells: cells are added in pairs as above. Shrinking the interval $[a, b]$ if necessary, so that only one such pair (α, β) is present, we see that all other cells in the boundary of β were already present in $\mathcal{K}(a)$, and all cells containing α in their boundary have been assigned values greater than b (since f is a discrete Morse function), so that there is a deformation retract that leads from $\mathcal{K}(b)$ to $\mathcal{K}(a)$. This sort of deformation retract, where a pair of non-critical cells cancels out to give a space that is homotopy equivalent, is called an *elementary collapse*. Thus we conclude that the non-critical cells do not contribute to homology.

On the other hand, when $f^{-1}(a, b]$ contains a critical p -cell, as is the case for $f^{-1}(3, 4]$, then $\mathcal{K}(b)$ is homotopy equivalent to $\mathcal{K}(a)$ with a p -cell attached, a Fig. 3.4 illustrates. This is because, since the cell is critical, all cells in its boundary are already present in $\mathcal{K}(a)$, and all cells containing it are outside $\mathcal{K}(b)$, so that moving from $\mathcal{K}(a)$ to $\mathcal{K}(b)$ consists precisely of attaching this single cell.

The discussion so far indicates that when non-critical cells can be paired subject to certain conditions, then that pair can safely be removed from the complex. This is the idea we use to construct a discrete Morse function on an arbitrary regular complex.

One consequence of Lemma 3.1 is that, for any cell α , exactly one of the following holds:

- $s(\alpha) = 1$, in which case α lies in exactly one such pair

- $i(\alpha) = 1$, in which case α lies in exactly one such pair
- $s(\alpha) = i(\alpha) = 0$, in which case α is critical.

It appears that just describing the pairs that do not contribute to homology might be enough for our purposes, and that is indeed true: we do not need the actual values assigned to the cells, but only the pairs they are a part of.

Definition 3.4. A **discrete vector field** V on \mathcal{K} is a collection of pairs $(\alpha^{(p)}, \beta^{(p+1)})$ where $\alpha < \beta$, such that each cell is in at most one pair of V .

Suppose there is a discrete Morse function $f: \mathcal{K} \rightarrow \mathbb{R}$. The three alternatives listed above ensure that if we pair the cells $\alpha^{(p)} < \beta^{(p+1)}$ whenever $f(\alpha) \geq f(\beta)$, the resulting collection of pairs is a discrete vector field, called the **gradient vector field** of f .

Is every discrete vector field the gradient vector field of some discrete Morse function? If so, then we can do away with the function itself and just work with the vector field, for what we really need is just the pairs. However, that is not the case. A discrete vector field just requires that each cell be in at most one pair, so the collection

$$\left\{ \left(\alpha_1^{(p)}, \beta_1^{(p+1)} \right), \left(\alpha_2^{(p)}, \beta_2^{(p+1)} \right) \right\}$$

is a vector field. However, if $\alpha_1, \alpha_2 < \beta_1, \beta_2$, then it is not the gradient vector field of any discrete Morse function f , since we would require that such an f satisfy

$$f(\alpha_1) \geq f(\beta_1) > f(\alpha_2) \geq f(\beta_2) > f(\alpha_1), \quad (3.2)$$

which is not possible. This motivates the following definition:

Definition 3.5. Given a discrete vector field V on \mathcal{K} , a **V -path** is a sequence of cells

$$\alpha_0^{(p)}, \beta_0^{(p+1)}, \alpha_1^{(p)}, \beta_1^{(p+1)}, \dots, \beta_r^{(p+1)}, \alpha_{r+1}^{(p)} \quad (3.3)$$

such that, for each $0 \leq i \leq r$,

- $\alpha_{i+1} \neq \alpha_i$,
- (α_i, β_i) is a pair in V and
- $\beta_i > \alpha_{i+1}$

We say such path is a **closed path** if $\alpha_0 = \alpha_{r+1}$.

As in Eq. (3.2), if Eq. (3.3) is a path of a gradient vector field, it must satisfy the condition

$$f(\alpha_0) \geq f(\beta_0) > f(\alpha_1) \geq f(\beta_1) > \dots \geq f(\beta_r) > f(\alpha_{r+1}). \quad (3.4)$$

Clearly, this can happen only if the path is not closed. We therefore arrive at the following result:

Proposition 3.3. *If V is the gradient vector field of some discrete Morse function, then there are no closed V -paths.*

It turns out that the converse is also true, that is, the existence of closed paths is the only obstruction to a discrete vector field being a gradient vector field:

Theorem 3.4. [For02, Theorem 3.5] *A discrete vector field is the gradient vector field of a discrete Morse function if and only if there are no closed V -paths.*

We will not prove the remaining direction here. It follows from the following theorem of graph theory:

Theorem 3.5. [For02, Theorem 3.6] *Let G be a directed graph. Then there is a real-valued function of the vertices that is strictly decreasing along each directed path if and only if there are no directed loops.*

The Hasse diagram of any poset can be modeled as a directed graph thus: (here we consider the face poset of \mathcal{K}) draw a directed edge from $\beta^{(p+1)}$ to $\alpha^{(p)}$ if and only if $\alpha < \beta$. So, for the cell decomposition shown in Fig. 3.1, the directed graph (hereafter called only *graph*) is as in Fig. 3.5.




Figure 3.5: The Hasse diagram of the decomposition in Fig. 3.1 modeled as a directed graph.

Now, consider the gradient vector field of the function f shown in Fig. 3.3. The pairs for this vector field are $(3,2,1)$ & $(23,1)$; $(1,3,2)$ & $(13,2)$; and $(2,1,3)$ & $(12,3)$. We indicate this in the graph by reversing the corresponding edges. Fig. 3.6 shows this.




Figure 3.6: The directed graph for the discrete Morse function f .

Theorem 3.5 says that, once we do this for a discrete vector field, the resultant graph has no directed loops if and only if it is the gradient vector field of a discrete Morse function. A directed loop can be described as: a way to start from a node and keep following outbound edges until you return to the node. As the reader can check, Fig. 3.6 has no directed loops, as one would expect from a gradient vector field.

We now incorporate some useful terminology from graph theory (where our graphs are always finite):

Definition 3.6. A **matching** M on a graph G is a set of edges such that no two vertices of G share an edge in M .

Definition 3.7. A matching on (the Hasse diagram of) the face poset of a regular cell complex is **acyclic** if the directed graph obtained by directing matching edges upward and all other edges downward has no directed cycles.

Thus a discrete Morse function on a regular cell complex is nothing but an acyclic matching on its face poset.

Definition 3.8. A matching M on a graph G is called **maximum** (respectively, **maximum acyclic**) if it contains the largest number of edges of any matching (respectively, acyclic matching) of G . A matching M is called **maximal** (respectively, **maximal acyclic**) if, when a new edge is added to M it no longer remains a matching (respectively, acyclic matching).

A graph may have multiple maximum matchings, but they must all have the same cardinality. A maximum matching is always maximal, but the converse is not true.

When can we say that the graph has a maximum acyclic matching? One answer is provided by the weak Morse inequalities, which hold true for discrete Morse functions.

For a discrete Morse function f , denote by m_i the number of critical cells of dimension i . For a field \mathbb{F} , denote by β_i the i^{th} \mathbb{F} -Betti number of \mathcal{K} , that is,

$$\beta_i = \dim H_i(\mathcal{K}; \mathbb{F}).$$

Theorem 3.6. [For98] (The weak Morse inequalities) For each $0 \leq i \leq n$, where n is the highest dimension of any cell of \mathcal{K} ,

$$m_i \geq \beta_i.$$

Definition 3.9. Fix a field \mathbb{F} . A discrete Morse function f for which the weak Morse inequalities become equalities is called an **\mathbb{F} -perfect Morse function**.

The weak Morse inequalities immediately give the following result:

Corollary 3.7. If a Morse function is perfect, then its gradient vector field is a maximum acyclic matching.

Remark 3.1. There is a notion, in graph theory, of a *perfect matching*. This is not to be confused with our notion of a perfect Morse function. We will have no occasion to use perfect matchings in this thesis.

3.2 A perfect Morse function on \mathcal{K}_L^*

This section follows closely the construction described in [PZ15] by G. Panina and A. Zhukova. We fix a generic length vector $L = (l_1, \dots, l_n)$ and assume, as before, that $l_1 \leq \dots \leq l_n$. We now describe the construction of a discrete Morse function on the cell complex \mathcal{K}_L^* , which is dual to the regular cell complex \mathcal{K}_L on \mathcal{M}_L . Recall that the cells are labeled by cyclically ordered admissible partitions and we may assume that the n -set occurs last.

Since the homology groups of \mathcal{M}_L are torsion-free (Theorem 1.4), we can extend the definition of a perfect Morse function to when the coefficients take values in \mathbb{Z} . The Morse inequalities still hold, and we call \mathbb{Z} -perfect functions simply *perfect Morse functions*. The discrete Morse function that we now describe is, in this sense, perfect.

In what follows, we will often omit the phrase “is labeled by” and simply equate a cell with its label. The following **notation** and definition will be heavily used in this and subsequent chapters:

- The ellipsis “...” denotes any (possibly empty) ordered admissible collection of subsets of $[n]$ (that is, each subset in the collection is short).
- The asterisk “*” denotes any (possibly empty) subset of $[n]$. By abuse of notation, it also denotes the elements of such a set. For example, $\{3, *\}$ is the same as $\{3\} \cup *$.
- The letter N denotes the n -set (that is, the set containing n).
- The triangle \blacktriangledown , with or without a subscript, denotes (a possibly empty) string of singletons arranged in decreasing order. For example, \blacktriangledown_1 could represent the sequence $\{4\}\{2\}\{1\}$ but not the sequence $\{4\}\{1\}\{2\}$ or the sequence $\{4\}\{1, 2\}$.
- The expression “ $k < I$ ” indicates that $k < i$ for each $i \in I$. Similarly, the expression “ $k < \blacktriangledown$ ” indicates that k is less than the element in each singleton of \blacktriangledown .

Definition 3.10. For $k \in [n]$, a set $I \subset [n]$ is called **k-prelong** if I itself is short, but $I \cup \{k\}$ is long.

We can now proceed with the construction, which is divided into several steps:

Step 1: Pair together the cells

$$\alpha = (\dots \{1\} I \dots) \text{ and } \beta = (\dots \{1\} \cup I \dots) \quad (3.5)$$

if $n \notin I$ and $\{1\} \cup I$ is short (which must anyway hold true if the cell β is to exist at all).

Observe that the cell α lies in the boundary of β , so this is a valid pairing. Moreover, by construction, no cell can be a part of more than one such pairing. The only possible paths at this stage are of the form

$$\alpha_0, \beta_0, \alpha_1, \beta_1, \dots, \alpha_p, \beta_p, \alpha_{p+1}, \quad (3.6)$$

where each pair (α_i, β_i) is of the form described in Eq. (3.5).

Claim 3.8. *In the path described by Eq. (3.6) above, α_{p+1} cannot equal α_0 .*

Proof. If $p = 0$, then the claim is true by the definition of a path. If $p > 0$, then, by the same definition, we have that, for each $0 \leq i \leq p$, the cell α_i is distinct from α_{i+1} . Therefore, if α_i is of the form $(\dots \{1\} I \dots)$, then of necessity β_i must be $(\dots \{1\} \cup I \dots)$, and since α_{i+1} is not α_i , it must be $(\dots I \{1\} \dots)$ (since it has to be of the form described in Eq. (3.5)). Proceeding thus, we see that as we traverse the path, the element 1 keeps moving to the right, entering and exiting sets as necessary.

If α_{p+1} is to equal α_0 , then 1 must move into the n -set and reemerge on the left by a splitting of the n -set, but that is forbidden by the rules of Step 1. \square

This claim has shown that there are no cycles at this stage. In other words, the discrete vector field described by the pairing thus far is a discrete Morse function. We didn't need to know this right now (after all, what matters is that we have a discrete Morse function when *all* the steps are completed), but the foregoing proof illustrates the point of the condition $n \notin I$ that was imposed.

What are the unpaired cells at this stage? If an unpaired cell has 1 as a singleton, then the Step 1 conditions ensure that the set following $\{1\}$ must be either the n -set, or a 1-prelong set. If 1 is not a singleton, and if 1 and n are in separate sets, then the 1-set can be split to obtain a cell with 1 as a singleton followed by the rest of the 1-set, so the cell can be paired. We conclude that the unpaired cells are of the types:

1. $(\dots \{1\} \{n, *\})$,
2. $(\dots \{1\} (\text{a 1-prelong set}) \dots)$ and
3. $(\dots \{1, n, *\})$.

Clearly, any cells of the above types are also unpaired, so we have completely characterized the unpaired cells at this stage.

Step 2: Pair together the cells

$$\alpha = (\dots \{2\} I \dots) \text{ and } \beta = (\dots \{2\} \cup I \dots) \quad (3.7)$$

if $n \notin I$; $\{2\} \cup I$ is short, and α and β have not yet been paired.

Observe that the above conditions automatically ensure that $1 \notin I$. For if 1 were in I , then since we have that α and β are unpaired, and since $n \notin I$, we must have $I = \{1\}$ by the characterization at the end of Step 1. But then $\beta = (\dots \{1, 2\} \dots)$ would be paired with the cell $(\dots \{1\} \{2\} \dots)$ in Step 1 itself, a contradiction.

As with step 1, we can give here also a characterization of the unpaired cells, and show that the pairing process so far yields no closed paths. But there is not much to be gained from this extra effort, and we postpone these tasks to the end of the construction.

The construction now proceeds similarly for $n - 2$ steps, with the k^{th} step looking thus:

Step k: Pair together the cells

$$\alpha = (\dots \{k\} I \dots) \text{ and } \beta = (\dots \{k\} \cup I \dots) \quad (3.8)$$

if $n \notin I$; $\{k\} \cup I$ is short, and α and β have not yet been paired.

Analogously to the situation in Step 2, the conditions here ensure that all elements in I are greater than k . For if not, and I is a singleton, then the same argument as in Step 2 can be used to show that β is already paired; if I is not a singleton, then its least element, say j , can be ejected from it to yield new cells

$$\alpha' = (\dots \{k\} \{j\} I' \dots) \text{ and } \beta' = (\dots \{j\} \{k\} \cup I' \dots),$$

where $I' = I \setminus \{j\}$. Then it can be seen that (α, α') and (β, β') must have been paired no later than Step j .

We illustrate this pairing procedure with a simple example.

Example 3.2. Let $L = (1, 1, 1, 2)$. Then, as we have seen already, \mathcal{M}_L is a circle, and the face poset of \mathcal{K}_L^* looks thus (as before, we dispense with set braces and write elements of a single set contiguously; we further dispense with the 4-set because it is always singleton in this case):




Figure 3.7: The Hasse diagram of the dual face poset for $L = (1, 1, 1, 2)$

The pairing now proceeds as follows (recall that *any* two element subset of $[4]$ not containing the element 4 is short):

In Step 1, the following cells get paired:

- $(1, 2, 3)$ and $(12, 3)$
- $(2, 1, 3)$ and $(2, 13)$
- $(1, 3, 2)$ and $(13, 2)$
- $(3, 1, 2)$ and $(3, 12)$

In Step 2, the following cells get paired:

- $(2, 3, 1)$ and $(23, 1)$

Thus the only unpaired cells are $(1, 23)$ and $(3, 2, 1)$. As discussed in the previous section, this pairing can also be depicted as a matching on a directed graph, by making all edges in the face poset point downward, and then reversing the directions of the matched edges (Fig. 3.8).

As the reader can verify, there are no cycles in this graph, hence the pairing yields a discrete vector field which is in fact a discrete Morse function.




Figure 3.8: Matching on the directed graph for $L = (1, 1, 1, 2)$. The unpaired cells are enclosed in boxes.

Similar to the example, the pairing construction always gives a discrete Morse function:

Proposition 3.9. [PZ15, Proposition 3.1] (1) Assume we have a gradient path of the discrete vector field given by the pairing construction. Assume also that $m < k$, and a cell

$$\alpha = (\dots \{k, *\} \dots \{m, *\} \dots)$$

belongs to the path (that is, the elements k and m belong to different sets and the set containing k is somewhere to the left of the set containing m).

Then, during the path after the cell α , k always remains to the left of m . In other words, cells of the following types do not occur in the gradient path after α :

- $(\dots \{k, m, *\} \dots)$
- $(\dots \{m, *\} \dots \{k, *\} \dots)$

(2) The introduced discrete vector field is a discrete Morse function.

Proof. (1) For a cell of either of the two types listed above to occur in the gradient path after α , the cell of the first type must occur at some point, because the alternative—that m pass through the n -set and reemerge on left of k —is forbidden. However, the cell having m and k in the same set cannot occur, because for that to happen, k must enter the m -set, which is forbidden since $k > m$.

(2) Suppose we have a closed path

$$\alpha_0, \beta_0, \alpha_1, \beta_1, \dots, \alpha_p, \beta_p, \alpha_{p+1},$$

with $\alpha_{p+1} = \alpha_0$ and the α_i 's unique for each $0 \leq i \leq p$. Suppose

$$\alpha_0 = (\dots \{k\} I \dots) \text{ and } \beta_0 = (\dots \{k\} \cup I \dots).$$

For α_p (which by hypothesis is distinct from α_0), there are three possibilities:

1. If α_p gets paired at the k^{th} step, then

$$\alpha_p = (\dots \{k\} J I \dots) \text{ and } \beta_p = (\dots \{k\} \cup J I \dots),$$

in which case we would have

$$\alpha_0 = (\dots J \{k\} I \dots).$$

2. If α_p gets paired at the j^{th} step for $j > k$, then

$$\alpha_p = (\dots \{k\}K \dots \{j\}J \dots) \text{ and } \beta_p = (\dots \{k\}K \dots \{j\} \cup J \dots),$$

where K is k -prelong, because this is the only way in which α_p can remain unpaired after the k^{th} step. In this case we have

$$\alpha_0 = (\dots \{k\}I I' \dots),$$

where $K = I \cup I'$.

3. If α_p gets paired at the j^{th} step for $j < k$, then there are only two ways in which α_0 can remain unpaired after the j^{th} step:

- (a) $\alpha_p = (\dots \{k\}I \dots \{j\}JK \dots)$ and $\beta_p = (\dots \{k\}I \dots \{j\} \cup JK \dots)$, where K is j -prelong, in which case we have $\alpha_0 = (\dots \{k\}I \dots J \{j\}K \dots)$.
- (b) $\alpha_p = (\dots \{k\}I \dots \{j\}J \{n, *\})$ and $\beta_p = (\dots \{k\}I \dots \{j\} \cup J \{n, *\})$, in which case we have $\alpha_0 = (\dots \{k\}I \dots J \{j\} \{n, *\})$.

In case 1, there exists some $m \in J$ such that $m > k$. Then, m is to the left of k in α_0 but to the right in α_p , which is impossible by part (1). In case 2, there has to be some step at which the set K gets formed. But at that step k can still be moved into some set, or the set containing k can be split, so the cell would have been paired at the k^{th} step itself. So this scenario is also not possible.

In case 3(a), there exists $m > j$ in J , but then m and j are not allowed to cross each other by part (1). Similarly, case 3(b) is also not possible. \square

Another feature of Example 3.2 which holds in the general case is the type of unpaired cells one is left with. Recall that the unpaired cells in the example were $(3, 2, 1, 4)$ and $(1, 23, 4)$: one having a sequence of singletons going in the decreasing order, and the other having a singleton which could have moved into the set following it but for the fact that it would have created a long subset. The following theorem states that, in essence, these are the only types of unpaired cells possible. Using the terminology of the previous section, we may call the unpaired cells *critical cells*.

Theorem 3.10 (Theorem 4.1). [PZ15] *The critical cells of the discrete Morse function described so far are exactly cells of the following two types:*

- Cells of **type 1** labeled by

$$(\blacktriangledown \{n, *\})$$

- Cells of **type 2** labeled by

$$(\blacktriangledown_1 \{k\} I \blacktriangledown_2 \{n, *\})$$

where the following hold:

1. I is a k -prelong set,

2. $k < I$ and
3. $k < \blacktriangledown_1$.

Proof. For any singleton in \blacktriangledown in a cell of type 1, the set following it either contains an element less than it, or contains n . So no singleton can be moved into a subsequent set to form a pair. Moreover, in a pairing, the n -set cannot split, and that is the only non-singleton set. Hence a cell of type 1 cannot be paired. The same reasons apply to cells of type 2 as well. In addition, since I is k -prelong, k cannot move into I to give a pairing. So we only need to check that I cannot split to give a pairing, but that is impossible since the resulting cell would be (where $I' = I \setminus \{j\}$):

$$(\blacktriangledown_1\{k\}\{j\}I'\blacktriangledown_2\{n,*\}),$$

which would be paired not with the cell we began with, but with

$$(\blacktriangledown_1\{k,j\}I'\blacktriangledown_2\{n,*\}).$$

We have therefore shown that a cell of either of the above two types must be critical.

Conversely, suppose α is a critical cell. If α contains only singletons (other than the n -set), then they must either all be in decreasing order (to prevent any two from merging), ensuring α is of type 1, or there exists not more than one pair of singletons in the ascending order whose union is a long subset. (Such a situation is indeed possible: consider, for $L = (1, 2, 2, 2)$, the cell $(2, 3, 1, 4)$ is critical because $\{2, 3\}$ is a long subset.) In the latter case, α is of type 2.

On the other hand, if α contains a non-singleton set I other than the n -set, it must be preceded by a singleton $\{k\}$ such that $k < I$, for otherwise I itself could split to give a pairing of

$$(\dots\{j\}I'\dots) \text{ and } \alpha = (\dots I \dots),$$

where j is the least element of I and $I' = I \setminus \{j\}$. Moreover, I must be k -prelong, else k could enter I to give a pairing of α with the resultant cell. If we can now show that α can contain not more than two non-singleton sets, it would follow that α is of type 2. So suppose α contains L , a non-singleton set which is distinct from both I and the n -set. The same conditions that apply to I must also apply to L , that is, L is preceded by a singleton $\{k'\}$ such that $L' := L \cup \{k'\}$ is long. But then its complement is a short subset that contains $I \cup \{k\}$, a long subset, which is absurd. \square

Example 3.3. We can use this method to find a discrete Morse function on a sphere. Suppose $(L = 1, 1, \dots, 1, n - 1 - \epsilon)$. As we know (Example 1.8), in this case \mathcal{M}_L is the $(n - 3)$ -sphere. According to the theorem, the only critical cells of the discrete Morse function on this sphere are:

- Type 1: $(\{n - 1\} \dots \{2\}\{1\}\{n\})$, since the only short subset containing n is the singleton.
- Type 2: $(\{1\}\{2, 3, \dots, n - 1\}\{n\})$, since the only long subset not containing n is $[n - 1]$.

Thus we have one critical 0-cell and one critical $(n - 3)$ -cell. Since the sphere has Betti number 1 precisely in dimensions 0 and $n - 3$, we conclude that we in fact have a perfect Morse function.

In general, however the method above does not yield a perfect Morse function:

Example 3.4. Suppose $L = (1, 1, 1, 1, 1)$. Then, as we know (Theorem 1.9), \mathcal{M}_L is diffeomorphic to the orientable surface of genus 4. The following are all critical 0-cells of type 2:

- $(\{1\}\{2,3\}\{4,5\})$
- $(\{1\}\{2,4\}\{3,5\})$
- $(\{1\}\{3,4\}\{2,5\})$
- $(\{2\}\{3,4\}\{1,5\})$

Clearly, this Morse function is not perfect, and in fact, is quite far from being perfect.

We therefore need some way to reduce the number of critical cells if we wish to have a perfect Morse function. This is done by a technique wherein certain paths between cells are identified and then reversed. It is based on the following theorem, which is a discrete analog of the ‘‘Cancellation Theorem’’ in classical Morse theory:

Theorem 3.11. [*For02, Theorem 9.1*] Suppose f is a discrete Morse function on \mathcal{K} such that $\beta^{(p+1)}$ and $\alpha^{(p)}$ are critical, and there is exactly one gradient path from β to α . Then there is another Morse function g on \mathcal{K} with the same critical cells as f except that α and β are no longer critical. Moreover, the gradient vector field associated to g is equal to the gradient vector field associated to f except along the unique gradient path from β to α .

Proof. The hypotheses of the theorem imply that the directed graph of the Hasse diagram of \mathcal{K} has an acyclic matching, with a unique path between $\beta^{(p+1)}$ and $\alpha^{(p)}$, say

$$\beta, \alpha_0, \beta_0, \dots, \alpha_m, \beta_m, \alpha.$$

Then we ‘‘reverse’’ this matching, that is, we match α with β_m , α_m with β_{m-1} and so on, until we match α_0 with β . Since the path between β and α was unique, no cycles are introduced, and now α and β are no longer unmatched. \square

The technique described in the proof is known as **path reversal**. We wish to apply this technique to our current situation, but first we need a pair of critical cells that have a unique gradient path between them.

Proposition 3.12. *Suppose*

$$\beta = (\blacktriangledown_1\{k\}I\blacktriangledown_2\{n, j, *\}) \text{ and } \alpha = (\blacktriangledown_1\{k\}I\blacktriangledown_2 \cup \{j\}\{n, *\})$$

are critical cells of type 2 (where $\blacktriangledown_2 \cup \{j\}$ indicates that the set $\{j\}$ is appended to \blacktriangledown_2 and then moved to the appropriate position to ensure that \blacktriangledown_2 remains a descending sequence of singletons).

If I is j -prelong, then the cells are connected by exactly one gradient path.

Proof. Clearly there is at least one path: j gets ejected (to the left) from the n -set and keeps moving till α is reached. For another path to exist, I would have to split in β (for j cannot get ejected to the right of the n -set: I being j -prelong would prevent it from passing through). If that were the case, the least element of I , say i would have to enter I at some point in order to obtain α . At that stage, k cannot be to the right of I (for then it can never re-emerge to the left), nor can it be to the left of I (for then it would merge with i). Thus no such path exists. \square

Thus we have a pair of critical cells between which there is exactly one path. We seek to reverse this path. However, since it is possible to create cycles when reversing more than one path simultaneously, we impose some extra conditions.

Path Reversal Step: Reverse the path between the critical cells

$$\beta = (\nabla_1\{k\}I \nabla_2\{n, j, *\}) \text{ and } \alpha = (\nabla_1\{k\}I \{j\} \nabla_2\{n, *\})$$

if the following hold:

1. $j > *$,
2. $j > \nabla_2$ and
3. $j > k$.

These conditions ensure that we end up with a discrete vector field. Note that here we do not impose the condition that I be j -prelong; the condition $j > k$ is enough to ensure that j cannot get ejected to the right of the n -set if it is to reappear to the right of k .

Thus we have a discrete vector field, with a reduced number of unpaired cells. As the next results state, this is in fact a perfect Morse function, and so, in our situation at least, path reversal does the job.

Proposition 3.13 (Proposition 6.3). *[PZ15] The path reversal technique yields a discrete Morse function.*

Theorem 3.14 (Theorem 6.4). *[PZ15] The resulting discrete Morse function is perfect.*

Thus the proposition asserts that the final matching is acyclic, and the theorem that it is a maximum matching.

We also describe here the **critical cells** that remain after the path reversal, because we will need these later:

1. All cells of type 1, and
2. All cells $(\nabla_1\{k\}I \nabla_2\{n, *\})$ of type 2 (that is, $k < I$, $k < \nabla_1$ and I is k -prelong) such that, in addition,

$$k > * \text{ and } k > \nabla_2.$$

Chapter 4

The \mathbb{Z}_2 -action on \mathcal{M}_L

When \mathcal{M}_L is a manifold, there are no collinear configurations, so reflecting a polygon P about the X -axis (which is an involution) yields a polygon Q that is necessarily distinct from P . In other words, the group \mathbb{Z}_2 acts freely on \mathcal{M}_L . The orbit space (or quotient) of this action, denoted by \mathcal{O}_L , is therefore a manifold, with $\pi : \mathcal{M}_L \rightarrow \mathcal{O}_L$ a covering map. This chapter is devoted to the quotient manifold \mathcal{O}_L . The main questions we seek to answer are: does \mathcal{O}_L admit a cell structure with a convenient combinatorial description, and, if so, can we profitably do discrete Morse theory on it?

4.1 \mathbb{Z}_2 -equivariance

The first question can be answered almost immediately. We start by introducing some notions from [Str11] about group actions on a CW complex.

Definition 4.1. A map $X \rightarrow Y$ of CW complexes is called a **cellular map** if it restricts to a map of k -skeleta $X^k \rightarrow Y^k$ for each k . In particular, a cellular map takes k -cells to k -cells.

Definition 4.2. An action of a group G on a CW complex X is a **cellular action** (or G **acts cellularly**) if $g : X \rightarrow X$ is a cellular map for each $g \in G$.

A cellular map that fixes a cell might not fix each point of the cell. For example, consider the CW structure on the closed interval $[0, 1]$ consisting of two 0-cells and one 1-cell (with obvious attaching maps). Reflection about the point $\frac{1}{2}$ is a cellular map that fixes the 1-cell, but not the points of the 1-cell.

Definition 4.3. Given a group G acting on a CW complex X , X is called a **G -CW complex** if G acts cellularly and if, whenever $g \in G$ fixes a cell σ , it also fixes all points of σ .

The following proposition is a classical result of homotopy theory:

Proposition 4.1. [Geo07, Proposition 3.2.2] *Let X be a G -CW complex and let $\pi : X \rightarrow X/G$ be the quotient map onto the orbit space. Then X/G admits a CW structure whose cells are $\{\pi(e) \mid e \text{ is a cell of } X\}$.*

The next lemma will allow us to use the above results.

Lemma 4.2. \mathcal{M}_L is a \mathbb{Z}_2 -CW complex.

Proof. Since the only non-trivial element of \mathbb{Z}_2 acts by reflection (which does not fix any cell), we only have to check that reflection, which we denote by r , is a cellular map. Suppose P is a convex polygon with the angles of its sides satisfying $\theta_1(P) < \dots < \theta_n(P)$. Reflection takes it to a polygon $Q = r(P)$ with $\theta_i(Q) = -\theta_i(P)$ for all $1 \leq i \leq n-1$, and $\theta_n(Q) = \theta_n(P)$. Thus the angles of Q satisfy $\theta_{n-1}(Q) < \dots < \theta_1(Q) < \theta_n(Q)$.

Now suppose P is any polygon without parallel edges, with angles satisfying $\theta_{\lambda(1)}(P) < \dots < \theta_{\lambda(n)}(P)$, for some permutation $\lambda \in S_n$, which is in fact the label of the $(n-3)$ -cell P belongs to. Then reflection takes P to the polygon $Q = r(P)$ whose angles satisfy $\theta_{\lambda(n-1)}(Q) < \dots < \theta_{\lambda(1)}(Q) < \theta_{\lambda(n)}(Q)$.

Generalizing this to all polygons, we see that r maps a polygon with label $(I_1 \dots I_k)$, where $I_j \subset [n]$ and $n \in I_k$, to a polygon with label $(I_{k-1} \dots I_1 I_k)$. Since this is true for all polygons with that label, it is true for the cell itself, hence r maps cells to cells. \square

The description of the reflection r in terms of cell labels (enunciated in the preceding proof) is an important one, and we reiterate it here (since we are deliberately blurring the distinction between a cell and its label, we can also think of r as an involution on the set of admissible partitions of $[n]$):

$$\begin{aligned} r : \mathcal{M}_L &\rightarrow \mathcal{M}_L \\ (I_1 \dots I_k) &\mapsto (I_{k-1} \dots I_1 I_k). \end{aligned}$$

Combining Proposition 4.1 and Lemma 4.2, we get

Theorem 4.3. \mathcal{O}_L has a CW structure whose k -cells are labeled by equivalence classes of admissible partitions of $[n]$ into $k+3$ blocks, where two such partitions λ and λ' are equivalent if $\lambda' = r(\lambda)$. The cell (labeled by the equivalence class) $\bar{\lambda} = \overline{\lambda'}$ is contained in the boundary of $\bar{\sigma}$ if and only if λ is refined by either σ or $r(\sigma)$.

The last statement of the theorem makes sense because if λ refines σ , then $r(\lambda)$ refines $r(\sigma)$. In fact r induces a poset isomorphism of \mathcal{K}_L with itself. We denote the CW complex on \mathcal{O}_L described in Theorem 4.3 by \mathcal{C}_L , and its dual complex by \mathcal{C}_L^* . Some examples will illustrate the theorem.

Example 4.1. Consider our running example $L = (1, 1, 1, 2)$ with $\mathcal{M}_L = S^1$. The reflection r maps the cell $(\{1\}\{2,3\}\{4\})$ to the cell $(\{2,3\}\{1\}\{4\})$. The equivalence classes of 0-cells (with braces removed) are:

1. $\overline{(1, 23, 4)} = \overline{(23, 1, 4)}$
2. $\overline{(2, 13, 4)} = \overline{(13, 2, 4)}$
3. $\overline{(3, 12, 4)} = \overline{(12, 3, 4)}$

The equivalence classes of 1-cells are:

1. $\overline{(1, 2, 3, 4)} = \overline{(3, 2, 1, 4)}$
2. $\overline{(2, 1, 3, 4)} = \overline{(3, 1, 2, 4)}$
3. $\overline{(1, 3, 2, 4)} = \overline{(2, 3, 1, 4)}$

So $\mathcal{O}_L = S^1/\mathbb{Z}_2 = \mathbb{R}P^1 \simeq S^1$ has the cell structure shown in Fig. 4.1a.




Figure 4.1: The cell structure for \mathcal{O}_L is shown here, for two different values of L

Example 4.2. Let $L = (1, 2, 2, 2)$. Here \mathcal{M}_L is a disjoint union of two circles. Reflection maps one circle to the other. In Fig. 2.9, therefore, if one imagines a vertical line between the two circles, each cell is mapped to its image across this line. The quotient is a circle and its cell structure is shown in Fig. 4.1b.

The case $n = 4$ is not special: in general, when \mathcal{M}_L is a sphere, \mathcal{O}_L is the real projective space of the same dimension, and when \mathcal{M}_L is a disjoint union of two tori, \mathcal{O}_L is a torus.

4.2 The \mathbb{Z}_2 -homology of \mathcal{O}_L

To answer the second question posed at the beginning of this chapter, namely, whether discrete Morse theory can be done on the quotient space, we need to first understand what a perfect Morse function on this space is. For \mathcal{M}_L , Theorem 1.4 assures absence of torsion in the \mathbb{Z} -homology groups. For \mathcal{O}_L this is no longer true. For example, taking $L = (1, 1, 1, 1, 3)$, the space \mathcal{O}_L is $\mathbb{R}P^2$, whose \mathbb{Z} -homology groups are \mathbb{Z} , \mathbb{Z}_2 and 0 in dimensions 0, 1 and 2 respectively. So the \mathbb{Z} -ranks of $H_1(\mathbb{R}P^2)$ and $H_2(\mathbb{R}P^2)$ are 0, while we know that there is no cellular decomposition of $\mathbb{R}P^2$ that can be obtained without 1-cells or 2-cells. Hence there cannot be a \mathbb{Z} -perfect Morse function on \mathcal{O}_L .

Therefore, for the rest of this chapter, instead of considering \mathbb{Z} -perfect Morse functions, we look at \mathbb{Z}_2 -perfect Morse functions (that is, discrete Morse functions with as many critical k -cells as the \mathbb{Z}_2 -rank of the k th \mathbb{Z}_2 -homology group). Henceforth when we say *perfect Morse function*,

this is what we mean. Observe that the \mathbb{Z}_2 -homology groups of $\mathbb{R}P^2$ are \mathbb{Z}_2 in each dimension, so there is some possibility of finding a perfect Morse function.

In order to show that a particular discrete Morse function is perfect (or not), we first need to know the \mathbb{Z}_2 -Betti numbers of \mathcal{O}_L . We use a couple of results from the paper [HK98], by J.-C. Hausmann and A. Knutson, to derive what we need.

In [HK98], the polygon spaces $\text{Pol}(L)$ and $\text{Pol}_{\mathbb{R}}(L)$ are defined for generic L , the former being the space of polygons in \mathbb{R}^3 modulo $SO(3)$, the latter being the space of polygons in \mathbb{R}^2 modulo $O(2)$ (and therefore \mathcal{O}_L for us). It is shown that

Proposition 4.4. [HK98, Corollary 4.3] *The Poincaré polynomial of $\text{Pol}(L)$ is given by*

$$P_{\text{Pol}(L)}(t) = \frac{1}{1-t^2} \sum_J (t^{2(|J|-1)} - t^{2(n-|J|-1)}) \quad (4.1)$$

where the sum ranges over all short subsets J containing n .

Proposition 4.5. [HK98, Theorem 9.1] *As \mathbb{Z}_2 -vector spaces,*

$$\dim(H^k(\mathcal{O}_L; \mathbb{Z}_2)) = \dim(H^{2k}(\text{Pol}(L); \mathbb{Z}_2)) \quad (4.2)$$

Therefore, using Eq. (4.1), Eq. (4.2) and Poincaré duality, the “ \mathbb{Z}_2 -Poincaré polynomial” of \mathcal{O}_L (that is, a polynomial with coefficients the \mathbb{Z}_2 -dimensions of the homology groups), is given by

$$\tilde{P}_{\mathcal{O}_L}(t) = \frac{1}{1-t} \sum_J (t^{|J|-1} - t^{n-|J|-1})$$

Now suppose, for $0 \leq k \leq n-3$, the number of short subsets of cardinality $k+1$ containing n is denoted by a_k . Then we can rewrite the above expression as

$$\tilde{P}_{\mathcal{O}_L}(t) = \frac{1}{1-t} \sum_{k=0}^{n-3} a_k (t^k - t^{n-k-2}) \quad (4.3)$$

If $2k+2 < n$, then we have

$$\begin{aligned} \frac{t^k - t^{n-k-2}}{1-t} &= \frac{t^k}{1-t} (1 - t^{n-2k-2}) \\ &= t^k (1 + t + t^2 + \dots + t^{n-2k-3}) \\ &= t^k + t^{k+1} + \dots + t^{n-k-3}. \end{aligned}$$

If $2k+2 > n$, then we have

$$\begin{aligned} \frac{t^k - t^{n-k-2}}{1-t} &= -\frac{t^{n-k-2}}{1-t} (1 - t^{2k+2-n}) \\ &= -t^{n-k-2} (1 + t + \dots + t^{2k+1-n}) \\ &= -(t^{n-k-2} + t^{n-k-1} + \dots + t^{k-1}). \end{aligned}$$

So,

$$\tilde{P}_{\mathcal{O}_L}(t) = \sum_{k=0}^{\lfloor \frac{n-2}{2} \rfloor} a_k (t^k + \dots + t^{n-3-k}) - \sum_{k=\lfloor \frac{n}{2} \rfloor}^{n-3} a_k (t^{n-k-2} + \dots + t^{k-1}). \quad (4.4)$$

Given $0 \leq j \leq n$, the \mathbb{Z}_2 -rank of the j^{th} homology group has a contribution of a_k for each k such that $k \leq j \leq n - k - 3$, and a contribution of $-a_k$ for each k such that $n - k - 2 \leq j \leq k - 1$. That is, a_k contributes if and only if $k \leq j$ and $k \leq n - j - 3$, and $-a_k$ contributes if and only if $k \geq j + 1$ and $k \geq n - j - 2$. Define, for $0 \leq j \leq n - 3$,

$$\begin{aligned} p_j &:= \min(j, n - j - 3) \\ q_j &:= \max(j + 1, n - j - 2) \end{aligned}$$

Note that $p_j < \frac{n}{2} - 1 < q_j$. Also $p_j = p_{n-3-j}$ and $q_j = q_{n-3-j}$. Let

$$\beta_j(L) := \dim(H_j(\mathcal{O}_L; \mathbb{Z}_2)). \quad (4.5)$$

The preceding discussion now gives us

Theorem 4.6. *The \mathbb{Z}_2 -dimension of the j^{th} homology group of \mathcal{O}_L is*

$$\beta_j(L) = \sum_{k=0}^{p_j} a_k - \sum_{k=q_j}^{n-3} a_k \quad (4.6)$$

for all $0 \leq j \leq n - 3$.

We immediately note that $\beta_0(L) = \beta_{n-3}(L) = a_0 = 1$, so \mathcal{O}_L is always connected and (obviously) \mathbb{Z}_2 -orientable. Some examples will help us understand the computations involved in Eq. (4.6).

Example 4.3. Consider the length vector L for which \mathcal{M}_L is S^{n-3} , so the only short subset containing n is the singleton, that is, $a_0 = 1$ and $a_i = 0$ for $i > 0$. Then the above expression simplifies to

$$\beta_j(L) = a_0 = 1$$

for all $0 \leq j \leq n - 3$, which agrees with the result that \mathcal{O}_L is $\mathbb{R}P^{n-3}$.

Example 4.4. Assume that \mathcal{M}_L is not connected, so that $\{n - 2, n - 1\}$ is a long subset. Then $a_k = \binom{n-3}{k} = a_{n-3-k}$, as we know. Fix a j . If $j \leq n - j - 3$, then we also have $j + 1 \leq n - j - 2$, so $p_j = j$ and $q_j = n - j - 2$, and Eq. (4.6) becomes

$$\beta_j(L) = \sum_{k=0}^j a_k - \sum_{k=n-j-2}^{n-3} a_k.$$

Now if $k < j$, then $n - 3 - k > n - 3 - j$, so a_k gets canceled out by $-a_{n-3-k}$, and similarly if $k > n - j - 3$, so we get that

$$\beta_j(L) = a_j = \binom{n-3}{j}.$$

A similar argument can be used to get the same value for $j > n - j - 3$, and therefore this agrees with our result that \mathcal{O}_L is the $(n - 3)$ -torus.

Example 4.5. Assume $n > 4$, so that $q_j > 1$ for all j . Consider the length vector L for which \mathcal{M}_L is the product $S^1 \times S^{n-4}$, so the only short subsets containing n are the singleton and $\{1, n\}$ (Theorem 1.6), that is, $a_0 = 1$, $a_1 = 1$ and $a_i = 0$ for $i > 1$. Then (5) gives us

$$\beta_j(L) = \begin{cases} a_0 = 1, & j = 0, n-3, \\ a_0 + a_1 = 2, & \text{otherwise.} \end{cases}$$

4.3 Discrete Morse theory on \mathcal{C}_L^*

In this section, we answer the second question posed at the beginning of this chapter. We describe a discrete Morse function on the CW complex \mathcal{C}_L^* which is dual to the cell structure described in Section 4.1. The k -cells of this complex are labeled by equivalence classes of admissible partitions of $[n]$ into $n - k$ blocks. Before we proceed, we point out an important feature of the cells in \mathcal{K}_L (the complex for \mathcal{M}_L):

Definition 4.4. Let $\lambda = (I_1 \dots I_k)$ be a cyclically ordered partition of $[n]$ (which, if admissible, is a cell in \mathcal{K}_L). Let j be the greatest element outside the n -set (which is always assumed to be I_k). Suppose $j \in I_l$ for some $1 \leq l < k$. Let i be the greatest element outside the j -set and the n -set (in particular, $i \notin I_l, I_k$). Suppose $i \in I_m$ for some m such that $1 \leq m < k$ and $m \neq l$.

Then λ is said to be of **class** (\mathbf{i}, \mathbf{j}) if $m < l$ (that is, I_m is to the left of I_l), and of **class** (\mathbf{j}, \mathbf{i}) otherwise. The class of λ is denoted by $\text{cl}(\lambda)$.

Definition 4.5. A cell of class (i, j) is said to be **ascending** if $i < j$ and **descending** otherwise.

Definition 4.6. If $\text{cl}(\lambda) = (i, j)$ and $\text{cl}(\lambda') = (i', j')$ are two cells with $i \leq i'$ and $j \leq j'$, then λ' is said to be **higher** than λ and λ is said to be **lower** than λ' .

Some examples will help explain these definitions.

Example 4.6. Let $n = 7$. We omit the braces for the sets in the partition.

$\alpha = (23, 156, 47)$ is of class $(3, 6)$ (since 6 is the greatest element outside the 7-set, and 3 the greatest element outside both the 7-set and the 6-set). Since $3 < 6$, it is ascending.

$\beta = (156, 23, 47)$ is of class $(6, 3)$ and descending.

$\gamma = (12, 34, 567)$ is of class $(2, 4)$ and ascending.

$\delta = (13, 2, 4567)$ is of class $(3, 2)$ and descending.

α is higher than both γ and δ . δ is lower than β . Of α and β , neither is higher or lower than the other. Similarly, of γ and δ , neither is higher or lower than the other.

Observe in Example 4.6 that $\beta = r(\alpha)$, and that $\text{cl}(\beta)$ is just $\text{cl}(\alpha)$ flipped. This is true in general: reflection maps a cell of class (i, j) to one of class (j, i) . In particular, it maps an ascending cell to a descending cell and a descending cell to an ascending one. In fact, we have

Lemma 4.7. *The cells of \mathcal{K}_L are partitioned into two sets: one with the ascending cells, and one with the descending. Reflection establishes a bijection between these two sets:*

$$\{\text{Ascending cells}\} \xleftrightarrow{r} \{\text{Descending cells}\}$$

Each cell in the quotient complex \mathcal{C}_L is an equivalence class containing an ascending cell and a descending cell (each the reflection of the other). So in \mathcal{C}_L it does not make sense to talk of ascending or descending cells. But we can slightly modify the notion of class:

Definition 4.7. A cell $\bar{\lambda}$ in \mathcal{C}_L is said to be of **class** $\{i, j\}$ if one of its preimages under the quotient map (a cell in \mathcal{K}_L) is of class (i, j) (equivalently, of class (j, i)). As before, we denote the class by $\text{cl}(\bar{\lambda})$.

Clearly, the class $\{i, j\}$ is equal to the class $\{j, i\}$. We also have to modify the notion of “higher”:

Definition 4.8. If $\text{cl}(\bar{\lambda}) = \{i, j\}$ and $\text{cl}(\bar{\lambda}') = \{i', j'\}$ are two cells with $\min(i, j) \leq \min(i', j')$ and $\max(i, j) \leq \max(i', j')$, then $\bar{\lambda}'$ is said to be **higher** than $\bar{\lambda}$ and $\bar{\lambda}$ is said to be **lower** than $\bar{\lambda}'$.

We are now ready to begin our description of the discrete Morse function on \mathcal{C}_L^* (on which too the notions we have defined above hold). The construction is a modified version of that described in Section 3.2 and, in fact, begins as a discrete Morse function on \mathcal{K}_L^* which is then “pushed down” to \mathcal{C}_L^* . Thus we first outline, stepwise, the pairing that happens in \mathcal{K}_L^* , which is hereafter called the *modified matching*.

Step 1: Pair together the cells

$$\alpha = (\dots \{1\} I \dots) \text{ and } \beta = (\dots \{1\} \cup I \dots) \quad (4.7)$$

in \mathcal{K}_L^* if the following conditions hold:

1. $n \notin I$,
2. $\{1\} \cup I$ is short,
3. α is ascending and
4. $\text{cl}(\alpha) = \text{cl}(\beta)$.

Observe that 1) this is the same as Step 1 of the earlier pairing, with two extra conditions, and 2) conditions 3 and 4 together imply that β must also be ascending.

The rest of the steps proceed similarly, with the two extra conditions imposed at each step:

Step k: Pair together the cells

$$\alpha = (\dots \{k\} I \dots) \text{ and } \beta = (\dots \{k\} \cup I \dots) \quad (4.8)$$

in \mathcal{K}_L^* if the following conditions hold:

1. $n \notin I$,
2. $\{k\} \cup I$ is short,

3. α and β have not yet been paired,
4. α is ascending and
5. $\text{cl}(\alpha) = \text{cl}(\beta)$.

After the $(n - 2)$ -th step, we have the

Final step: If α and β have been paired in \mathcal{K}_L^* , then pair together $\bar{\alpha}$ and $\bar{\beta}$ in \mathcal{C}_L^* (where $\bar{\alpha}$ denotes the image of α under the map $\pi : \mathcal{K}_L^* \rightarrow \mathcal{C}_L^*$.)

Since we have paired only ascending cells in \mathcal{K}_L^* , any cell in \mathcal{C}_L^* can have at most one paired preimage, so the final step indeed gives a matching on the face poset of the quotient. We claim that this matching is acyclic, hence the pairing describes a discrete Morse function.

Lemma 4.8. *If $\bar{\alpha}$ is contained in the boundary of $\bar{\beta}$, then $\bar{\alpha}$ is higher than $\bar{\beta}$.*

Proof. Suppose $\text{cl}(\bar{\beta}) = \{i, j\}$ with $i > j$. Then, since each block of $\bar{\alpha}$ is a subset of a block of $\bar{\beta}$ (we are thinking of cells as admissible partitions now)—note that this is happening in the dual complex—the largest element outside the n -set in $\bar{\alpha}$ has to be greater than or equal to i , and similarly the largest element outside the n - and i -sets has to be greater than or equal to j . \square

Lemma 4.9. *If there is a gradient path*

$$\bar{\beta}_0, \bar{\alpha}_1, \bar{\beta}_1, \dots, \bar{\alpha}_p, \tag{4.9}$$

then $\bar{\alpha}_p$ is higher than $\bar{\beta}_0$.

Proof. Since we only pair cells of the same class, $\text{cl}(\bar{\alpha}_i) = \text{cl}(\bar{\beta}_i)$ for each $1 \leq i \leq p - 1$. Moreover, since $\bar{\alpha}_{i+1}$ is contained in the boundary of $\bar{\beta}_i$ for each i , it is higher than the latter by the previous lemma. The result follows. \square

Theorem 4.10. *The pairing on \mathcal{C}_L^* , as described in the final step above, gives a discrete Morse function.*

Proof. We need to show that the matching we have described is acyclic. So suppose there is a path

$$\bar{\alpha}_0, \bar{\beta}_0, \bar{\alpha}_1, \bar{\beta}_1, \dots, \bar{\alpha}_p \tag{4.10}$$

with $p > 1$ and $\bar{\alpha}_0 = \bar{\alpha}_p$. Since $\bar{\alpha}_0$ and $\bar{\beta}_0$ are paired, they are of the same class. By the previous lemma, each cell in the path after $\bar{\beta}_0$ is higher than $\bar{\beta}_0$. But $\bar{\alpha}_0 = \bar{\alpha}_p$, so, in fact, $\text{cl}(\bar{\alpha}_0) = \text{cl}(\bar{\beta}_0) = \text{cl}(\bar{\alpha}_1) = \dots = \text{cl}(\bar{\beta}_{p-1}) = \text{cl}(\bar{\alpha}_p)$.

We now “lift” this cycle to \mathcal{K}_L^* . Let α_0 be the ascending cell such that $\pi(\alpha_0) = \bar{\alpha}_0$. Let β_0 the cell with which α_0 is paired (in particular, $\pi(\beta_0) = \bar{\beta}_0$). Next, suppose α_1 is ascending with $\pi(\alpha_1) = \bar{\alpha}_1$. Note that $\text{cl}(\alpha_0) = \text{cl}(\beta_0) = \text{cl}(\alpha_1) = (i, j)$ for some $i < j$. If α_1 is not in the boundary of β_0 , then it must be in the boundary of $r(\beta_0)$ (for otherwise $\bar{\alpha}_1$ would not be in the boundary of $\bar{\beta}_0$). But since $\text{cl}(r(\beta_0)) = (j, i)$, we have a cell of class (i, j) in the boundary of a

cell of class (j, i) , which is impossible. Hence α_1 is in the boundary of β_0 . Continuing thus, we obtain a path

$$\alpha_0, \beta_0, \alpha_1, \beta_1, \dots, \alpha_p \quad (4.11)$$

with α_i and β_i ascending for each i (and, in particular, $\alpha_0 = \alpha_p$). Thus the cycle (4.10) in \mathcal{C}_L^* lifts to the cycle (4.11) in \mathcal{K}_L^* . The matching on the ascending cells is, however, a subset of the matching of Section 3.2, and is acyclic by Proposition 3.9. Hence the cycle in (4.10) cannot exist. \square

The critical cells of this discrete Morse function correspond to the critical ascending cells of the modified matching on \mathcal{K}_L^* we have just described. The next theorem gives an explicit description

Theorem 4.11. *The critical cells of the discrete Morse function on \mathcal{C}_L^* are images, under π , of the following types of ascending cells in \mathcal{K}_L :*

1. Cells of **type 1Q**, labeled by

$$(\{i\}J \blacktriangledown \{n, *\})$$

where

- $\blacktriangledown < i < J$
- $\{i\} \cup J$ is short

2. Cells of **type 2Q**, labeled by

$$(\blacktriangledown_1 \{k\}I \blacktriangledown_2 \{n, *\})$$

where

- I is a k -prelong set,
- $k < I$ and
- $k < \blacktriangledown_1$.

Proof. A cell of type 1Q is of class (i, j) , where j is the greatest element of J . In the usual matching, it would be paired with the cell $(\{i\} \cup J \blacktriangledown \{n, *\})$, which is a descending cell. Therefore, in the modified matching, a cell of type 1Q remains unpaired. A cell of type 2Q remains unpaired even in the original matching, hence is a critical cell even for the modified matching.

Conversely, suppose an ascending cell α^p of type (i, j) is unpaired. There can only be two possibilities: 1) α is a critical cell even for the original matching, and 2) the original matching would have paired α with a cell of a different class. If 1) holds, then α is of one of the two types described in Theorem 3.10. However, a critical cell of type 1 can never be ascending, hence α must be of type 2 (which is the same as type 2Q above). If 2) holds, then there are two possibilities:

1. α^p would have been paired with β^{p+1} :

The only way the class of α can change is if the pairing involves i entering J . Hence $\alpha = (\dots \{i\}J \dots)$ with $j \in J$ and $\{i\} \cup J$ short. However, since α remained unpaired until

the i^{th} step, the elements less than i and outside the n -set must be in singletons and arranged in descending order. Thus

$$\alpha = (\{i\}J \blacktriangledown \{n, *\})$$

and is of type 1Q. Observe that β is descending.

2. α^p would have been paired with γ^{p-1} :

If γ were ascending, then using the argument just concluded, we would have α descending. Hence γ is descending. Suppose $\gamma = (\dots R \dots S \dots)$ is of class (r, s) with $s < r$. But then, in any pairing involving a singleton in γ joining a set, the resulting cell would still be of class (r, s) , hence descending. So such a situation cannot arise.

Thus we have exhausted all possibilities for α . □

Just as there were a lot of critical cells (far greater than the Betti numbers, that is) at the end of the pairing process (and before path reversal) in Section 3.2, so there are a lot of critical cells here. We single out three cases in which this is not true, that is, in which the Morse function we have described is already perfect.

Example 4.7. Suppose \mathcal{M}_L is the sphere S^{n-3} . The only short subset containing n is the singleton. So the only possible critical cell of type 2Q is $(\{1\}\{2, \dots, n-1\}\{n\})$, which is an $(n-3)$ -dimensional cell. The critical cells of type 1Q are of the form

$$(\{i\}J \blacktriangledown \{n\}).$$

Since $\blacktriangledown < i < J$, such a cell is completely determined by i . So there is exactly one such critical k -cell for each $1 \leq k \leq n-4$. Now \mathcal{O}_L is $\mathbb{R}P^n$, which has $\beta_k(L) = 1$ for all $1 \leq k \leq n-3$. Hence the Morse function is perfect.

Example 4.8. Assume $n > 4$. Suppose \mathcal{M}_L is $S^1 \times S^{n-4}$. Then the only short subsets containing n are the singleton and $\{1, n\}$. The only possible critical cells of type 2Q are $(\{2\}\{3, \dots, n-1\}\{1\}\{n\})$ and $(\{2\}\{3, \dots, n-1\}\{1, n\})$, the former of dimension $(n-4)$ and the latter, $(n-3)$.

There is one 0-dimensional cell of type 1Q: $(\{n-2\}\{n-1\}\{n-3\} \dots \{1\}\{n\})$.

There is one $(n-4)$ -dimensional cell of type 1Q: $(\{3\}\{4, \dots, n-1\}\{2\}\{1, n\})$.

For each $0 < k < n-4$, there are two k -dimensional cells of type 1Q: one with 1 in the n -set, and one without.

Thus we have two critical cells in each dimension between 0 and $(n-3)$, and one each in dimensions 0 and $(n-3)$. Looking at the calculation performed in Example 4.5, we see that the Morse function is perfect.

Example 4.9. Suppose \mathcal{M}_L is the disjoint union of two tori, so $\{n-2, n-1\}$ is a long subset. Then \mathcal{O}_L is a torus. Since $\{n-2, n-1\}$ is long, there are no critical cells of type 1Q. The only critical cells of type 2Q are of the form

$$(\{n-2\}\{n-1\} \blacktriangledown \{n, *\}).$$

In dimension k (which requires $(n-k)$ blocks), there are exactly $\binom{n-3}{n-3-k} = \binom{n-3}{k}$ such cells, which equals $\beta_k(L)$. So the Morse function is perfect.

In conclusion, we state that the previous example, in some sense, is the “ideal” scenario. The next few results explain what that means. Here we extensively use the fact that

$$\mathcal{M}_L \text{ is disconnected} \iff \mathcal{M}_L \simeq T^{n-3} \sqcup T^{n-3} \iff \{n-2, n-1\} \text{ is long.}$$

Proposition 4.12. \mathcal{M}_L is disconnected if and only if the boundary of every ascending cell in \mathcal{K}_L (equivalently, \mathcal{K}_L^*) contains only ascending cells.

Proof. Suppose \mathcal{M}_L is disconnected. Then any ascending cell α looks like

$$\alpha = (\dots \{n-2, *\} \dots \{n-1, *'\} \dots \{n, *''\})$$

since $n, n-1$ and $n-2$ can never be in the same set. So, if α' is in the boundary of α , then, since α refines α' , the blocks of α' are formed by taking unions of blocks of α . In such blocks too the $(n-2)$ -set must precede the $(n-1)$ -set, so α' is also ascending.

Conversely, if $\{n-2, n-1\}$ is short, then there exists a 0-cell

$$\alpha = (\{n-2, n-1, *\} I \{n, *'\})$$

which is descending and is contained in the boundary of the ascending cell

$$\beta = (\{n-2\} \{n-1, *\} I \{n, *'\}),$$

hence there is at least one ascending cell that contains a descending cell in its boundary. \square

Proposition 4.13. If \mathcal{M}_L is disconnected, then one connected component is made up entirely of ascending cells, and the other entirely of descending cells.

Proof. For the purposes of this proof, the term *walk* refers to a walk in the face poset of \mathcal{K}_L (considered as an undirected graph) in the graph-theoretic sense. That is,

$$\alpha_1, \alpha_2, \dots, \alpha_p$$

is a walk if, for each i , $\alpha_i < \alpha_{i+1}$ or $\alpha_i > \alpha_{i+1}$.

If \mathcal{M}_L is disconnected, so is its face poset. Hence, a walk can exist between two cells only if they are in the same component. The converse is also true: given two cells α and β , there exists a walk between α and some 0-cell α_0 (for instance, by repeated refinement), and similarly between β and some β_0 . It suffices to show that there is a walk between α_0 and β_0 . But these are points in the manifold, and are therefore connected by a path. Successively listing all cells this path encounters gives a walk between α_0 and β_0 .

Now suppose \mathcal{M}_L is disconnected. By the argument just concluded, it suffices to show that there is always a walk between two ascending cells, and never one between an ascending cell and a descending cell. By Proposition 4.12, every ascending cell contains an ascending 0-cell in its boundary, and (analogously) every descending cell contains a descending 0-cell. Let α_0 and β_0 be 0-cells such that α_0 is ascending. Appealing to the reasoning in the preceding paragraph, we need to show that:

- (1) If β_0 is ascending, then there is a walk between α_0 and β_0 .
 (2) If β_0 is descending, then there is no such walk.

For (1), let $k < n - 2$, and suppose

$$\alpha_0 = (\{n - 2, k, *\} \{n - 1, *'\} \{n, *''\}).$$

Then

$$\begin{aligned} \alpha_0 = & (\{n - 2, k, *\} \{n - 1, *'\} \{n, *''\}), \\ & (\{n - 2, *\} \{k\} \{n - 1, *'\} \{n, *''\}), \\ & (\{n - 2, *\} \{n - 1, k, *'\} \{n, *''\}), \\ & (\{n - 2, *\} \{n - 1, *'\} \{k\} \{n, *''\}), \\ & (\{n - 2, *\} \{n - 1, *'\} \{n, k, *''\}) \end{aligned}$$

is a walk that takes $k < n - 2$ to any other block in the partition. Thus any element less than $n - 2$ can be freely moved around. Since β_0 is also ascending, we can thus construct a walk from α_0 to β_0 by just moving the elements to their desired location, because at no point do any two out of $n - 2$, $n - 1$ and n ever end up in the same set.

For (2), suppose there is such a walk. But since β_0 is descending, there must be some cell in the walk in which $n - 2$ and $n - 1$ are in the same set, which is impossible. \square

Thus the ascending cells form one torus, the descending cells the other torus, and the \mathbb{Z}_2 -involution is a poset isomorphism (apart from also being a diffeomorphism) between the two.

Theorem 4.14. \mathcal{K}_L^* admits a \mathbb{Z}_2 -equivariant perfect Morse function if and only if \mathcal{M}_L is disconnected.

Proof. Suppose \mathcal{M}_L is disconnected. Then the modified discrete Morse function on the ascending cells has critical cells as described in Example 4.9: exactly $\binom{n-3}{k}$ ascending critical k -cells (apart from the descending cells; these are still unmatched). Using the poset isomorphism between the ascending cells and descending cells that we have obtained, this matching induces a discrete Morse function on the descending cells (that is, we match α and β if $r(\alpha)$ and $r(\beta)$ are matched). By definition, this Morse function is \mathbb{Z}_2 -equivariant. The number of critical cells is just twice that of just the ascending cells: $2\binom{n-3}{k}$ critical k -cells, which is the k^{th} Betti number of \mathcal{M}_L . Thus we have a \mathbb{Z}_2 -equivariant perfect Morse function on \mathcal{K}_L^* . Note that this function is different from the one described in Section 3.2: that one is also perfect, but not \mathbb{Z}_2 -equivariant.

Conversely, suppose we have a \mathbb{Z}_2 -equivariant perfect Morse function on \mathcal{K}_L^* . Then, if α is a critical 0-cell, so is $r(\alpha)$. But r fixes no cells, hence these two cells are distinct. Since the Morse function is perfect, $H_0(\mathcal{M}_L)$ has rank at least two, so \mathcal{M}_L is disconnected. \square

Chapter 5

Looking ahead

So far, we have described a regular cell structure on \mathcal{M}_L which has a combinatorial interpretation. We have used this combinatorial interpretation and discrete Morse theory to construct a CW complex which has exactly as many cells as the Betti numbers of \mathcal{M}_L . We have studied the orbit space $\mathcal{O}_L = \mathcal{M}_L/\mathbb{Z}_2$, which is a manifold because reflection about the X -axis (which generates the \mathbb{Z}_2 -action) is free at the level of cells as well as individual points. On this quotient manifold, we have described a cell structure whose existence is an immediate consequence of the \mathbb{Z}_2 -invariance of the cell structure on \mathcal{M}_L . Thus this cell structure also has a combinatorial interpretation, which we have exploited to do discrete Morse theory on it, and obtain a discrete Morse function which, in certain cases, is \mathbb{Z}_2 -perfect.

In the final chapter, we discuss some topics on which work is ongoing, and some for which there is scope for future work. Most of these arise as reasonable extensions of the results we have described so far.

5.1 A minimal complex for \mathcal{M}_L

The fact that the homology of \mathcal{M}_L is torsion-free (Theorem 1.4) was proved by Farber and Schütz in [FS07] using classical Morse and Morse-Smale theory. Theorem 3.14 uses this result to show that the Morse function in Section 3.2 is perfect. Consequently, the resulting CW complex is a minimal complex modeling the homotopy type of \mathcal{M}_L .

In light of the above discussion one can ask the question: can it be shown that the critical cells that remain at the end of the process in Section 3.2 form a CW complex for \mathcal{M}_L with trivial attaching maps? If this can be shown, then we would have an alternative proof of Theorem 1.4, because the Betti numbers would be precisely the number of critical cells. The resulting proof would therefore provide some combinatorial insights which the original proof doesn't. In this subsection we report on the progress we have made towards answering this question

In the discrete setting, the attaching maps of the CW complex are dependent only on the gradient paths between critical cells of successive dimension. So the homology can be computed as follows. Let $C_p(\mathcal{K}_L^*)$ denote the free abelian group generated by critical p -cells. The boundary

operator $\partial: C_{p+1}(\mathcal{K}_L^*) \rightarrow C_p(\mathcal{K}_L^*)$ is given by

$$\partial(\beta^{p+1}) = \sum_{\alpha^p} c_{\alpha,\beta} \cdot \alpha$$

where

$$c_{\alpha,\beta} = \sum_{\gamma \in \Gamma(\beta,\alpha)} m(\gamma),$$

$\Gamma(\beta, \alpha)$ being the set of gradient paths from the boundary of β to α . The multiplicity $m(\gamma)$ of a gradient path is ± 1 , depending on whether, given γ , the orientation on β induces the chosen orientation on α , or the opposite orientation. However, if we use \mathbb{Z}_2 -coefficients, the task is simplified: the boundary maps depend only on the parity of number of paths between any pair of critical cells.

We now present, in detail, a characterization and the number of gradient paths between critical cells β^{p+1} and α^p . The analysis is split into four separate cases:

Case 1: α and β are both of type 1.

Lemma 5.1. *Suppose $\beta = (\blacktriangledown\{n, *\})$ and $\alpha = (\blacktriangledown'\{n, *'\})$. Then there is a gradient path between β and α if and only if $*' \subset *$. If this condition holds, then there are exactly two gradient paths.*

Proof. Suppose $*'$ contains an element that $*$ doesn't. Then, during a gradient path from β to α , that element must enter the n -set at some point. But during a pairing, no element can enter the n -set. So there are no gradient paths. Conversely, suppose $*' \subset *$ holds. Then we must have (omitting braces for singletons and writing $*$ as $m, +$)

$$\beta = (i_1 i_2 \dots i_k \{n, m, +\}) \text{ and } \alpha = (i_1 i_2 \dots i_j m i_{j+1} \dots i_k \{n, +\}),$$

where $i_1 > i_2 > \dots > i_j > m > i_{j+1} > \dots > i_k$. Then there are two gradient paths: m gets ejected either forwards or backwards from the n -set and ends up between i_j and i_{j+1} after successful merges and splits. Clearly these are the only possible paths. \square

Case 2: β is of type 1 and α of type 2.

This case has been dealt with in [PZ15]:

Proposition 5.2. [PZ15, Proposition 5.2] *There are no gradient paths from a critical cell of type 1 to a critical cell of type 2.*

Case 3: β is of type 2 and α of type 1.

Lemma 5.3. *Suppose \mathcal{M}_L is connected. Let $\beta = (\blacktriangledown_1\{k\}I\blacktriangledown_2\{n, *\})$ be a critical cell with the usual inequalities. Then there is a gradient path from β to a critical cell $\alpha = (\blacktriangledown'\{n, *'\})$ if and only if I has two elements and $* = *'$. When these conditions hold, there are exactly two gradient paths and \blacktriangledown' is obtained by splitting up the elements of I , inserting them in the appropriate positions in \blacktriangledown_1 and then concatenating the result with k and \blacktriangledown_2 .*

Proof. Suppose there is a gradient path between β and α . Then, since α is just one dimension lower and must have only singletons outside the n -set, I cannot contain more than 2 elements. However, since $\{n-1, n-2\}$ is short (by assumption, \mathcal{M}_L is connected), I cannot contain just one element. Hence $|I| = 2$. Further, the n -set cannot now split during the gradient path because all join steps will create a non-singleton outside the n -set (hence at each split step that non-singleton needs to be split). So $* = *'$. Conversely, if these conditions hold, then there are two paths from β to α : split I , and then keep doing the usual pairing and splitting till you obtain \blacktriangledown' as described in the statement of the lemma. There are two paths because I can be split in two ways. \square

When \mathcal{M}_L is not connected, it is possible a similar result exists, though we haven't directed our energies towards it yet.

Case 4: α and β are both of type 2.

This is work in progress; we hope that here too we get an even number of gradient paths.

A summary of the number of possible gradient paths is shown in Table 5.1.

β is of type	α is of type	Number of paths
1	1	0 or 2
1	2	0
2	1	0 or 2
2	2	Unknown

Table 5.1: Possible number of gradient paths from β to α when \mathcal{M}_L is connected.

5.2 Topology of \mathcal{O}_L

In Chapter 4, we presented some new ideas: a combinatorially-described cell structure on the quotient \mathcal{O}_L , and a discrete Morse function on its dual complex. Central to the proofs that this discrete Morse function is perfect in some cases is Theorem 4.6, which tells us the precise ranks of the \mathbb{Z}_2 -homology groups. Thus Theorem 4.6 performs the same function for \mathcal{O}_L as Theorem 1.4 does for \mathcal{M}_L : it allows us to construct a minimal cell complex.

We ask, as we did in the previous section, whether it is possible to show that there is always an even number of gradient paths between cells of successive dimension in the dual complex of \mathcal{O}_L . If we can do this, then we would have an alternative proof of Theorem 4.6.

Before we can do that, however, we need to find a way to reduce the number of critical cells we end up with at the end of the matching process on \mathcal{C}_L^* in Section 4.3. We are investigating whether it is possible to employ a modified path reversal technique as described at the end of Section 3.2 in order to do this.

Other questions of a similar flavor include:

1. What are the integral Betti numbers of \mathcal{O}_L ? This can be answered using the *transfer homomorphism*; the covering map $\pi : \mathcal{M}_L \rightarrow \mathcal{O}_L$ induces an injective homomorphism in rational cohomology whose image is the invariant subgroup (see, for example, [Hat02, Proposition 3G.1]). Consequently, in our situation, we have

$$H^i(\mathcal{O}_L; \mathbb{Q}) \simeq H^i(\mathcal{M}_L; \mathbb{Q})^{\mathbb{Z}_2} \quad \forall i.$$

The complex \mathcal{K}_L^* respects the \mathbb{Z}_2 action and can be used to determine the invariant subgroup.

2. Is there a way to construct a CW complex on \mathcal{O}_L whose cells and boundary maps have combinatorial descriptions (as we have done), and which gives the \mathbb{Z} -homology?
3. In [HK98] the authors compute the cohomology ring of \mathcal{O}_L with \mathbb{Z}_2 -coefficients. What is this ring when one considers \mathbb{Z} -coefficients?
4. Is it possible to compute the cohomology ring of \mathcal{M}_L using the answers to the two questions above, or otherwise?

A different line of inquiry concerns the results at the end of the previous chapter: when \mathcal{M}_L has two connected components, one consists solely of ascending cells and the other solely of descending cells. Call an ascending cell *pure* if it contains only ascending cells in its boundary. Then the disconnected \mathcal{M}_L has all its ascending cells pure, and they form a connected component.

For the sphere, the situation is slightly muddled. The pure ascending cells form a connected subset of the sphere, as do the (analogously defined) pure descending cells, with the impure cells of both kinds forming the part in between that deformation retracts onto the sphere of one less dimension.

So the question is: what sort of subset do the pure ascending cells form in general? What are its topological properties? In short, what does the ascending part “look like”?

5.3 Other directions

Finally, we come to a slightly different kind of question. Consider the following proposition, whose proof is fairly straightforward:

Proposition 5.4. *If $L = (l_1, \dots, l_n)$ is a generic length vector such that \mathcal{M}_L is not empty and $l_1 \leq \dots \leq l_n$, then the collection of short subsets of $[n]$, denoted by \mathcal{S}_L , has the following properties:*

- (1) \mathcal{S}_L contains all singletons;
- (2) If a set A is in \mathcal{S}_L , so are all of A 's subsets (so \mathcal{S}_L is actually an abstract simplicial complex);
- (3) From every pair of complementary subsets of $[n]$, \mathcal{S}_L contains exactly one; and

(4) If a set $A = \{a_1, a_2, \dots, a_k\}$ is in \mathcal{S}_L so is any set $B = \{b_1, b_2, \dots, b_k\}$ where $b_i \leq a_i$ for all i .

Now, suppose there exists a collection \mathcal{S} of subsets of $[n]$. Let i_1 be an element of $[n]$ appearing in a maximal number of sets in \mathcal{S} . Rename it 1. In $[n] \setminus \{i_1\}$, let i_2 be an element appearing in a maximal number of sets in \mathcal{S} . Rename it 2. We continue this process till we reach a stage where $i < j$ if and only if j appears in fewer or an equal number of sets in \mathcal{S} as i . Then, the above four conditions are necessary for \mathcal{S} to be \mathcal{S}_L for some L . (Our guess is that they are sufficient too, but we will not discuss that here.)

Now suppose there is a collection of subsets \mathcal{S} which, after the rearrangement described in the previous paragraph, satisfies the first three conditions above, but not the fourth. Then \mathcal{S} cannot be \mathcal{S}_L for any L , but it still defines a CW complex by the construction in Chapter 2.

Question: Is this CW complex homotopy equivalent to a moduli space of *something*? If so, what?

Next, suppose we relax condition (3). Then what kind of complex do we end up with? If it has sufficiently nice properties, can we answer the above question?

Bibliography

- [Far08] M. Farber. *Invitation to Topological Robotics*. Zurich Lectures in Advanced Mathematics. European Mathematical Society, 2008.
- [FS07] M. Farber and D. Schütz. *Homology of planar polygon spaces*. D. Geom Dedicata. Vol. 125 (2007), 75–92. arXiv: [math/0609140](#) [math.AT].
- [For98] R. Forman. *Morse Theory for Cell Complexes*. Advances in Mathematics. Vol. 134 No. 900145 (1998), 90–145.
- [For02] R. Forman. *A User's Guide to Discrete Morse Theory*. Seminaire Lotharingin de Combinatoire. Vol. 48 (2002).
- [Geo07] R. Geoghegan. *Topological Methods in Group Theory*. Graduate Texts in Mathematics. Springer New York, 2007.
- [Hat02] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
- [HK98] J.-C. Hausmann and A. Knutson. *The cohomology ring of polygon spaces*. Annales de l'institut Fourier. Vol. 48 (1998), 281–321. arXiv: [dg-ga/9706003](#) [math.DG].
- [KM95] M. Kapovich and J. Millson. *On the moduli space of polygons in the Euclidean plane*. J. Differential Geom. Vol. 42 No. 1 (1995), 133–164.
- [Mil64] J. Milnor. *Differential Topology*. Lectures on Modern Mathematics, vol. II. Ed. by T. L. Saaty. New York: John Wiley & Sons Inc., 1964. Chap. 6, 165–183.
- [Pan12] G. Panina. *Moduli space of planar polygonal linkage: a combinatorial description*. ArXiv e-prints (2012). arXiv: [1209.3241](#) [math.AT].
- [PZ15] G. Panina and A. Zhukova. *Discrete Morse theory for moduli spaces of flexible polygons, or solitaire game on the circle*. ArXiv e-prints (2015). arXiv: [1504.05139](#) [math.AT].
- [Str11] J. Strom. *Modern Classical Homotopy Theory*. Graduate Studies in Mathematics. American Mathematical Society, 2011.
- [Zie95] G. M. Ziegler. *Lectures on Polytopes*. Graduate Texts in Mathematics. Springer New York, 1995.