PL MORSE THEORY

MLADEN BESTVINA

Every mathematician has a secret weapon.
Mine is Morse theory.
–Raoul Bott

Contents

1. Goals

Morse theory is an extremely versatile tool, useful in a variety of situations and parts of topology and geometry. In these introductory lectures we will cover the foundations and discuss some typical applications. We will start by reviewing smooth Morse theory, then giving the PL counterpart. The rest of the sections consist of applications. The proofs are fairly detailed in the beginning but get sketchier as we go along. The reader is invited to find new applications.

These notes are a slightly expanded version of lectures I gave in 2007/08 in Berkeley (MSRI-Evans seminar), Les Diablerets, Switzerland (IIIe Cycle romand de Mathématiques), Baltimore (Johns Hopkins JAMI Conference) and in Osijek, Croatia (Fourth Croatian Mathematical Congress). I thank the organizers for these invitations.

2. Quick review: Smooth Morse theory

An excellent reference is Milnor's beautiful book [Mil63].

Let $f : M \to \mathbb{R}$ be a smooth function defined on a manifold M. A point $p \in M$ is *critical* if $df_p = 0$. In local coordinates, p is critical if all partial derivatives vanish at p. A critical point p is *non-degenerate* if the Hessian, i.e. the matrix of second partials at p, has nonzero determinant.

Lemma 2.1 (Morse Lemma). *Let p be a non-degenerate critical point. In suitable local coordinates around $p = 0$ the function f has the form*

$$f(x_1, \cdots, x_n) = (const) - x_1^2 - x_2^2 - \cdots - x_\lambda^2 + x_{\lambda+1}^2 + \cdots + x_n^2$$

The number $\lambda = 0, 1, \cdots, n$ is the index of f at p and it is independent of the choice of the suitable local coordinates (and can be defined independently of local coordinates in terms of the signature of the Hessian – which can itself be defined in a coordinate-free fashion).

A consequence of the Morse Lemma is that each non-degenerate critical point has a neighborhood in which there are no other critical points.

A function $f : M \to \mathbb{R}$ is *Morse* if all of its critical points are non-degenerate. One usually imposes other reasonable restrictions on f such as

- The image of f is bounded below.
- f is proper, i.e. sets $f^{-1}([a, b])$ are compact.
- Distinct critical points map to distinct points.

The knowledge of where the critical points are and their indices provides information about the manifold.

Theorem 2.2. *Let $f : M \to [0, \infty)$ be a proper Morse function. Then M is homotopy equivalent to a cell complex with the set of cells in 1-1 correspondence with critical points of f, the dimension of the cell equals to the index of the corresponding critical point.*

There is an even more precise description that recovers – in principle – M up to diffeomorphism and not only up to homotopy.

An n-dimensional λ-handle is the ball $D^\lambda \times D^{n-\lambda}$. The *attaching region* is $\partial D^\lambda \times D^{n-\lambda}$. A λ-handle is thought of as a thickening of a λ-cell and the attaching region is a thickening of the boundary of the cell.

Theorem 2.3. *Let $f : M \to [0, \infty)$ be a proper Morse function. Then M can be obtained (from the empty set) by successively attaching handles. Each handle is attached to the boundary of the union of the previous handles along its attaching region. The handles are in 1-1 correspondence with the critical points of f, the index of a handle equals the index of the corresponding critical point.*

There are also relative versions. The theorems are proved by considering the relationship between $N_t = f^{-1}[a, t]$ and $N_{t+s} = f^{-1}[a, t + s]$ for $s > 0$. If the closure of the difference $\overline{N_{t+s} - N_t}$ contain no critical points then one shows that N_t is diffeomorphic to N_{t+s}. This is done by considering the gradient flow (tapered within N_t) and flowing N_t onto N_{t+s}. If this closure contain one critical point and it has index λ and maps to $(t, t + s)$, then N_{t+s} is diffeomorphic to N_t with a handle of index λ attached.

Another useful thing to remember is that when f is replaced by $-f$, critical points of index λ become critical of index $n - \lambda$. This can be used to prove Poincaré duality.

Below we picture the obligatory example of the height function on the torus. There are four critical points, one of index 0 (the minimum), one of index 2 (the maximum) and two of index 1 (saddle points).

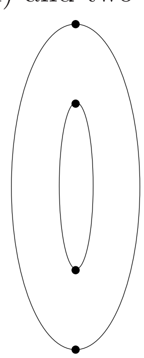

3. PL Morse theory

We want to develop a similar theory that applies to (nice) cell complexes instead of manifolds (see [BB97]). The applications we have in mind involve cube complexes and simplicial complexes. A natural class that includes both is the class of *affine polytope complexes*.

Definition 3.1. An APC is a cell complex X where each cell e is equipped with a characteristic function $\chi_e : C_e \to X$ with the following properties.

(1) Each C_e is a convex polyhedral cell in some $\mathbb{R}^m$ (fixed m for all cells).

(2) χ_e is an embedding for all e.
(3) The restriction of χ_e to any face of C_e agrees with the characteristic function of another cell precomposed with an affine homeomorphism of $\mathbb{R}^m$.

Replacing χ_e by precomposing it with an affine homeomorphism of $\mathbb{R}^m$ doesn't change anything of substance and we consider APC's related in this way as identical.

For example, finite dimensional simplicial complexes are APC's with each C_e a regular simplex (of appropriate dimension) in $\mathbb{R}^m$. Similarly, finite dimensional cube complexes are APC's.

One could easily allow infinite dimensional APC's by not insisting on all C_e's being contained in $\mathbb{R}^m$ for a fixed m, but to allow m to depend on C_e. Then in (3) we include the smaller Euclidean space in the larger in the standard way.

Definition 3.2. Let X be an APC and $f : X \to \mathbb{R}$ a function. We say that f is a *Morse function* provided

(1) **(affine)** For every cell e the composition $f\chi_e : C_e \to \mathbb{R}$ is the restriction of an affine function $\mathbb{R}^m \to \mathbb{R}$.
(2) **(no horizontal cells)** If $f|\ e$ is constant then $\dim\ e = 0$.
(3) The image $f(X^{(0)})$ of the 0-skeleton is a discrete subset of $\mathbb{R}$.

We can always position C_e in $\mathbb{R}^m$ (by precomposing with an affine homeomorphism) so that $f\chi_e$ is the restriction of the height function.

Examples 3.3. Let X be (the geometric realization of) a simplicial complex and $f : X \to \mathbb{R}$ a function. Condition (1) amounts to requiring that

$$f\big(\sum t_i v_i\big) = \sum t_i f(v_i)$$

for any point in a simplex with vertices v_i and barycentric coordinates t_i.

Another important class of examples is that of cube complexes. We will some applications later in the notes.

The next proposition claims that vertices are the only critical points.

Proposition 3.4. *Let $J = [a, b] \subset \mathbb{R}$ be a closed interval such that $(a, b]$ is disjoint from $f(X^{(0)})$. Then $f^{-1}(J)$ deformation retracts to $f^{-1}(a)$.*

Proof. Define $A_i = f^{-1}(a) \cup (f^{-1}(J) \cap X^{(i)})$. Thus

$$f^{-1}(a) = A_{-1} \subset A_0 \subset \cdots \subset A_m = f^{-1}(J)$$

and it suffices to argue that A_{i+1} deformation retracts to A_i. Consider an $(i+1)$-cell e and identify it with $C_e \subset \mathbb{R}^m$ via χ_e. Arrange things so that f is a height function on e. Then $e \cap f^{-1}(J)$ is a convex polytope. There are two horizontal faces, the bottom face that maps to a and the top face that maps to b (conceivably the bottom face degenerates to a vertex in the event that e contains a vertex that maps to a). Radial deformation from a point just above the top face gives a deformation retraction of $e \cap f^{-1}(J)$ to its boundary minus the top face. Putting together all these deformation retractions over all $(i+1)$-cells yields a deformation retraction of A_{i+1} to A_i. □

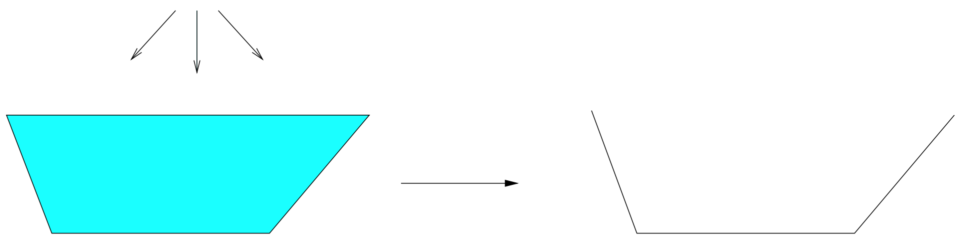

To understand the change in homotopy type at each vertex we introduce the notion of a *descending link*.

Definition 3.5. Fix a vertex $v \in X$. A cell e of X with $v \in e$ is *descending* if $f|e$ attains its maximum at v. The *descending link* at v, $Lk_{\downarrow}(v, X)$, is the link of v in the union of all descending cells.

Similarly, replacing "maximum" by "minimum", we have the notion of *ascending cells* and the *ascending link* $Lk_{\uparrow}(v, X)$.

Examples 3.6. In both examples below the Morse function is the height function and we consider the central vertex v.

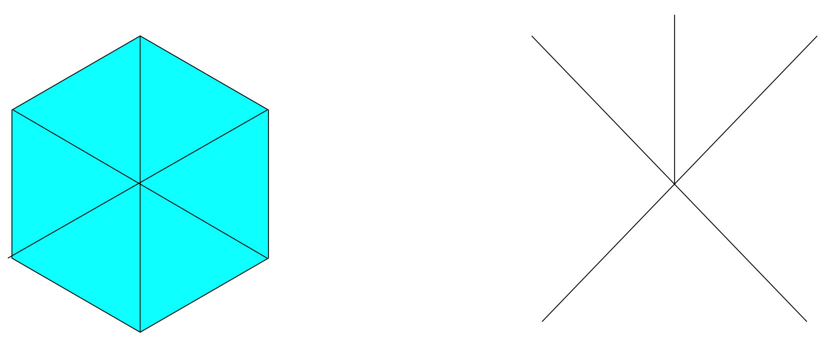

In the first example, both $Lk_{\downarrow}(v, X)$ and $Lk_{\uparrow}(v, X)$ are isomorphic to an arc triangulated with 3 vertices, since two of the triangles are descending, two are ascending, and two are neither. In the second example, $Lk_{\downarrow}(v, X)$ consists of two points, and $Lk_{\uparrow}(v, X)$ of three.

Proposition 3.7. *Suppose that $J = [a, b]$ is a closed interval and that $f^{-1}(J)$ contains one vertex v and $f(v) = b$. Then the pair $(f^{-1}(J), f^{-1}(a))$*

is homotopy equivalent rel $f^{-1}(a)$ *to* $(Q, f^{-1}(a))$ *where* Q *is* $f^{-1}(a)$ *with the cone on* $Lk_{\downarrow}(v, X)$ *attached.*

Proof. For each descending cell $e \ni v$ let $S_e = e \cap f^{-1}[a, b]$. Note that S_e is the cone on $Lk_{\downarrow}(v, e) = Lk(v, e)$ with v the cone point and the base contained in $f^{-1}(a)$. Define A_{-1} as the union of $f^{-1}(a)$ and all S_e's over all descending cells e. Thus A_{-1} is $f^{-1}(a)$ with the cone on $Lk_{\downarrow}(v, X)$ attached. Now define A_i as $A_{-1} \cup (f^{-1}(J) \cap X^{(i)})$. Thus

$$A_{-1} \subset A_0 \subset A_1 \subset \cdots \subset A_m = f^{-1}(J)$$

and A_{i+1} deformation retracts to A_i as before. □

A similar statement holds when there is more than one vertex in $f^{-1}(b)$; the cones on the descending links are attached in a pairwise disjoint fashion. In the example below, the complex X is a disk with the structure of a square complex with four 2-cells, $f^{-1}(J)$ is the region between the two horizontal lines; it deformation retracts to the subset indicated in black – the union of the lower horizontal line and the cones on the descending links for the three vertices in the upper horizontal line.

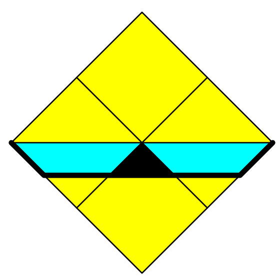

Analogously, if all vertices in $f^{-1}(J)$ are contained in $f^{-1}(a)$ then $f^{-1}(J)$ deformation retracts to $f^{-1}(b)$ with cones on the ascending links of vertices in $f^{-1}(a)$ attached. We will use these versions of Proposition 3.7 without further explanation.

4. Paths in a rectangle

Consider a $p \times q$ rectangular grid, such as the one pictured below. By A and B denote the lower left and upper right corners respectively.

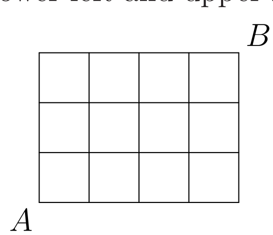

We want to organize the set of geodesic paths in the grid from A to B into a simplicial complex. A *geodesic path* is one that crosses precisely $p+q$ edges, equivalently, it always travels east and north. To that end, define a simplicial complex $X_{p,q}$ as follows. Its vertices are the vertices in the grid except for A and B. A collection of vertices spans a simplex iff there is a geodesic from A to B that passes through all of them. Thus the maximal simplices of $X_{p,q}$ correspond to the geodesic paths from A to B.

For example, $X_{1,1}$ has two points, while $X_{1,2}$ is an arc with four vertices. Likewise, $X_{0,q}$ is empty if $q = 1$ and is a $(q-2)$-simplex if $q > 1$.

We now claim that if $p > 1$ or $q > 1$ then $X_{p,q}$ is contractible.

The natural Morse function F assigns to a vertex (of $X_{p,q}$, which is a vertex of the grid) its distance to A (positioned at the origin): $(x, y) \mapsto x + y$, and then it is extended to all simplices in the affine fashion. There are no horizontal edges (distinct vertices along a geodesic path have different distances from A). If v is a vertex with coordinates (z, w) then

$$Lk_{\downarrow}(v, X_{p,q}) \cong X_{z,w}$$

which is contractible, by induction, if $z > 1$ or $w > 1$. The picture is now clear: the minimum of F is 1 and is attained at the vertices $(1, 0)$ and $(0, 1)$. At the next height 2, the vertex $(1, 1)$ has the descending link S^0, and it connects the two vertices at height 1. The other two vertices at height 2 are $(0, 2)$ and $(2, 0)$ and they, as well as all vertices at height > 2, have contractible descending links, so our space is contractible.

5. Configuration spaces

Let Y denote the tripod, pictured below, and let X_n be the configuration space of n distinct unmarked points in Y, that is,

$$X_n = \{(y_1, \cdots, y_n) \in Y^n | y_i \neq y_j \text{ for } i \neq j\}/S_n$$

where the symmetric group S_n acts by permuting the coordinates.

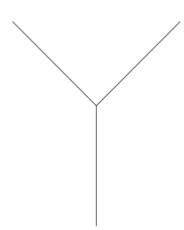

There is a natural function defined on X_n. Let $h : Y \to \mathbb{R}$ be the height function on Y (as it is drawn in the plane) and define

$$H : X_n \to \mathbb{R}$$

by

$$H(y_1, \cdots, y_n) = h(y_1) + \cdots + h(y_n)$$

A technical issue here is that X_n is not a cell complex, and if we equip X_n with a cell complex structure the image of the set of vertices will not be discrete. To overcome this problem, we will construct PL approximations to X_n. Fix a triangulation T of Y and define X_n^T as the subset of X_n consisting of equivalence classes of tuples $(y_1, \cdots, y_n)$ such that for $i \neq j$ the open simplices containing y_i and y_j have disjoint closures. Then X_n^T is naturally a cube complex and the restriction of H is a Morse function. The descending links are contractible except at the vertices of the form as pictured below. If there is at least one y_i on each arm then the descending link consists of two points (corresponding to moving the lowest points on the arms to the center point), and if there are no points on the arms then the descending link is empty. In all other cases the descending link is contractible. Assuming the triangulation T is fine enough, we now see that X_n^T is homotopy equivalent to the wedge of circles. E.g. for $n = 1$ the space is contractible, for $n = 2$ is homotopy equivalent to a circle, for $n = 3$ to the wedge of three circles etc. Moreover, this doesn't depend on the triangulation (as long as it is fine enough), and in fact X_n has that same homotopy type (this requires an argument, e.g. show that X_n deformation retracts to X_n^T).

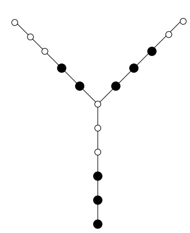

For more information about configuration spaces on graphs, see [FS05] (Morse theory used in that paper has slightly different flavor from the one here).

6. The Solomon-Tits theorem

Let K be a field.

6.1. **The Tits building for $SL_n(K)$.** The Tits spherical building X for $SL_n(K)$ is a simplicial complex with a vertex for each vector subspace of K^n of dimension $d = 1, 2, \cdots, n-1$ and a simplex for every flag, i.e. a chain of subspaces.

Theorem 6.1 (Solomon-Tits). *X has the homotopy type of the wedge of spheres of dimension $n-2$.*

Proof. Induction on n starting with $n = 2$ when $X = P^1K$, a discrete set.

Now suppose $n > 2$. Fix a line ℓ in K^n. Define

$$f : X^0 \to \{0, 1\} \times \{1, 2, \cdots, n-1\}$$

by

$$f(P) = (a, \dim P)$$

where $a = 0$ if $\ell \subset P$ and $a = 1$ if $\ell \not\subset P$. The target is ordered lexicographically. We could extend this to a function $X \to \mathbb{R}$ (after embedding $Im(f) \subset \mathbb{R}$ preserving the order) but we won't bother. The absolute minimum is $(0, 1)$ realized (only) on ℓ. Adjacent vertices map to distinct points. The descending links are as follows. Let $f(P) = (a, \dim P)$.

If $a = 0$ and $P \neq \ell$ then $Lk_\downarrow(P)$ is the complex of proper subspaces of P that contain ℓ; in particular this is a cone with cone point ℓ.

If $a = 1$ and $\dim P < n-1$ then $Lk_\downarrow(P)$ contains two kinds of vertices: proper subspaces of P as well as proper subspaces of K^n that contain $span(P \cup \ell)$. This is also a cone with cone point $span(P \cup \ell)$.

If $a = 1$ and $\dim P = n-1$ then $Lk_\downarrow(P)$ consists of proper subspaces of P, which is the Tits building for $SL_{n-1}(K)$.

Theorem now follows. □

Exercise 1. *If K is a finite field compute the number of spheres in the wedge.*

6.2. **The Tits building for $Sp_{2n}(K)$.** Now X is the flag complex of nontrivial isotropic subspaces in $K^{2n} = < a_1, b_1, \cdots, a_n, b_n >$ equipped with the symplectic form ω so that $\omega(a_i, a_j) = \omega(b_i, b_j) = 0$ and $\omega(a_i, b_j) = \delta_{ij}$.

Theorem 6.2. *X is homotopy equivalent to the wedge of $(n-1)$-spheres.*

Proof. Fix a line ℓ and define $f : X^0 \to \{0,1,2\} \times \{-(n-1), -(n-2), \cdots, -1, 1, 2, \cdots, n-1\}$ by

$$f(P) = \begin{cases} (0, \dim P) & \text{if } P \supseteq \ell \\ (1, \dim P) & \text{if } P \not\supseteq \ell, P \subseteq \ell^\perp \\ (2, -\dim P) & \text{if } P \not\subseteq \ell^\perp \end{cases}$$

Note that in the first case $P \subseteq \ell^\perp$ and in the last case $P \cap \ell^\perp$ has codimension 1 in P. The descending links $Lk_\downarrow(P)$ are as follows:

- $\emptyset$ if $P = \ell$,
- cone with cone point ℓ if $P \neq \ell$ and $P \supset \ell$,
- cone with cone point $span(P \cup \ell)$ if P is as in case 2,
- cone with cone point $P \cap \ell^\perp$ if P is as in case 3 but $\dim P > 1$,
- full link, if P is a line that intersects $\ell^\perp$ trivially.

In the last case a vertex in $Lk(P)$ has the form $span(P \cup Q)$ where Q is a unique isotropic subspace of $span(P \cup \ell)^\perp \cong K^{2n-2}$, so inductively $Lk(P)$ is homotopy equivalent to the wedge of $(n-2)$-spheres. □

7. Finiteness properties of groups

An Eilenberg-MacLane space $K(\Gamma, 1)$ associated with a discrete group Γ is a cell complex X with $\pi_1(X) = \Gamma$ and with contractible universal cover. Any two $K(\Gamma, 1)$'s are homotopy equivalent, so e.g. (co)homology groups are invariants of Γ. The basic questions one asks in this theory (see [Bro94]) are the following:

- (Finiteness properties of Γ) Can one choose a $K(\Gamma, 1)$ with finite k-skeleton? If so, Γ is said to be of type F_k. E.g. F_1 is "finitely generated" and F_2 is "finitely presented". Can one choose it to be a finite complex? (This is type F.)
- (Geometric dimension) What is the smallest dimension of a $K(\Gamma, 1)$? This number is very closely related to the cohomological dimension of Γ. It is known that $cd(\Gamma) \leq gd(\Gamma)$ with equality in all cases except possibly for $cd = 2$ and $gd = 3$.
- (Bieri-Eckmann duality) If X is a finite $K(\Gamma, 1)$, is compactly supported cohomology $H^*_c(X; \mathbb{Z})$ concentrated in one dimension? Is it free abelian? If so, a more general version of Poincaré duality holds for Γ, called Bieri-Eckmann duality.

These questions can be attacked using Morse theory. For example, to investigate finiteness properties of Γ, one would do the following:

- find a "nice" $K(\Gamma, 1)$,
- find a nice Morse function $f : K(\Gamma, 1) \to \mathbb{R}$,

- compute the descending links and hope that only finitely many fail to be $(k-1)$-connected.

Example 7.1. Consider the homomorphism $f : F_2(a, b) \to \mathbb{Z}$ given by $f(a) = f(b) = 1$ (generator of $\mathbb{Z}$). We will prove using Morse theory that $H = Ker(f)$ is free of infinite rank.

Let T be the Cayley graph of F_2 (i.e. the universal cover of the wedge of two circles). It is a 4-valent tree whose vertices are labeled by elements of F_2. We then have $f : T^{(0)} = F_2 \to \mathbb{Z} \subset \mathbb{R}$ and we observe that f sends the endpoints of each edge to adjacent integers. Thus there is a natural extension of f to $F : T \to \mathbb{R}$ that sends every edge by an affine homeomorphism to a closed interval of length 1. It is useful to think of F as the height function. Every deck transformation $\alpha : T \to T$ shifts the heights of all points by the same amount equal to $f(\alpha)$. More formally,

$$F(\alpha(x)) = f(\alpha) + F(x)$$

(when x is a vertex this just states that f is a homomorphism). In particular, H is the subgroup of the deck group consisting of level-preserving transformations.

For every vertex $v \in T$ there are 4 edges containing it. Two of them are descending and two are ascending. So both the descending and the ascending link at v consist of two points.

Now let $X = T/H$. Then X is a $K(H, 1)$ and we study its homotopy type by analyzing a Morse function on X. The function $F : T \to \mathbb{R}$ descends to $\Phi : X \to \mathbb{R}$. The ascending and descending links are 0-spheres. Also, $\Phi : X^{(0)} \to \mathbb{Z}$ is a bijection. The situation is pictured below.

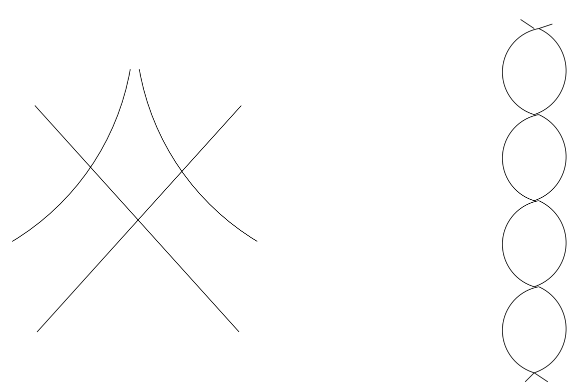

Of course, T doesn't want to be pictured in Euclidean plane with F corresponding to the height function, so pictured is only a very small part of T. We have much easier time drawing X, which is a chain of

$\mathbb{Z}$ circles, but note that for the Morse argument we only need to know "local information", i.e.

- that for each integer n the set $\Phi^{-1}(n)$ consists of exactly one vertex (which follows from the fact that $F^{-1}(n)$ is a single H-orbit of vertices, and
- that the descending and the ascending links of vertices in X consist of two points (which follows from the corresponding fact for F).

(it is the task of Morse theory to tell us about the topology of X; this example is so simple that we understand X without much of anything).

We now deduce from Proposition 3.7 (by starting with $\Phi^{-1}(0) = \{v_0\}$ and inductively understanding $\Phi^{-1}[-n, n]$) that X is homotopy equivalent to the wedge of circles with the set of circles naturally corresponding to $\mathbb{Z} - \{0\}$ and thus $H = \pi_1(X)$ is free of infinite rank.

Can we find a basis of H by looking at Φ?

Consider $X_n = \Phi^{-1}[-n, n]$. Thus $\pi_1(X_n) = \pi_1(X_{n-1}) * F_2$ and to find a basis for F_2 we need two loops α_n and α_{-n} in X_n based at v_0. The loop α_n will visit v_n once and otherwise it will stay in $\Phi^{-1}[-n, n)$, and in fact near v_n it will traverse the two descending cells at v_n. Such a loop can be constructed by finding a path in T that rises monotonically from the base vertex at height 0 to a vertex at height n, then descending back to height 0, using both descending edges at the highest point. For example $a^n b^{-n}$ will do.

The discussion for α_{-n} is similar. We find a path that descends from the base vertex to height $-n$ and rises back up to level 0, e.g. $a^{-n} b^n$.

Thus a basis for H is $\{a^n b^{-n} | n = \pm 1, \pm 2, \cdots\}$.

Example 7.2. Now take $f : F_2(a, b) \times F_2(x, y) \to \mathbb{Z}$, $f(a) = f(b) = f(x) = f(y) = 1$, and $H = Ker(f)$. For $K(F_2, 1)$ we take the wedge of two circles, and we put

$$K(F_2 \times F_2) = K(F_2, 1) \times K(F_2, 1)$$

so that it is the union of four tori that are glued to each other in a cyclic fashion, alternately along meridian and longitude.

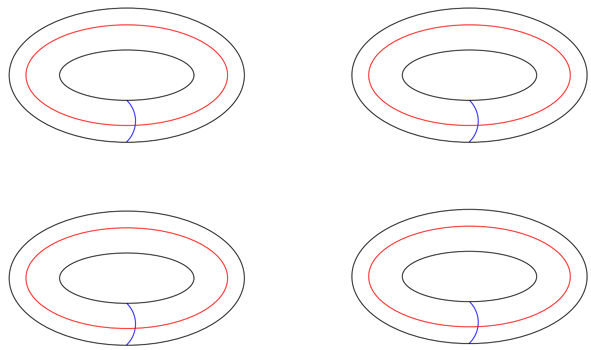

For $K(H,1)$ we take the infinite cyclic cover of $K(F_2 \times F_2, 1)$. More precisely, let

$$F : K(F_2 \times F_2, 1) \to S^1$$

be the map that is "addition" $S^1 \times S^1 \to S^1$ on each of the four tori (so it induces f in fundamental groups). Then we have the commutative diagram

$$\begin{array}{ccc} K(H,1) & \xrightarrow{\tilde{F}} & \mathbb{R} \\ \downarrow & & \downarrow \\ K(F_2 \times F_2, 1) & \xrightarrow{F} & S^1 \end{array}$$

Each torus in $K(F_2 \times F_2, 1)$ lifts to a cylinder in $K(H,1)$, which is therefore the union of four cylinders glued in a cyclic fashion along lines that represent lifts alternately of meridians and longitudes.

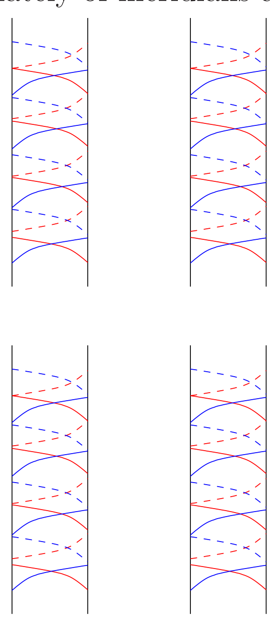

Now we need a Morse function on $K(F_2 \times F_2, 1)$, and we have a natural one, namely $\tilde{F}$.

At each vertex v of $K(H,1)$ both the descending and the ascending links are circles triangulated with four vertices (each cylinder contributes an arc).

The set $\tilde{F}^{-1}(0)$ is the wedge of four circles representing

$$ax^{-1}, ay^{-1}, by^{-1}, bx^{-1}$$

Homotopically, $K(H,1)$ is obtained from this graph by attaching infinitely many 2-cells. We deduce that H is generated by the four elements listed above and $H_2(H)$ is not finitely generated. It follows that H is not finitely presented.

Example 7.3. Let $f : F_2^n \to \mathbb{Z}$ send all $2n$ generators to 1, and let $H = Ker(f)$. Then $X = T^n/H$ has the homotopy type of a finite complex with infinitely many n-cells attached. The ascending and descending links are $(n-1)$-spheres. For $n = 3$ this is the Stallings' example of a finitely presented group whose H_3 is not finitely generated [Sta63]. For $n > 3$ these examples were first considered by Bieri [Bie76].

Example 7.4. The above examples can be generalized further, to the situation where the descending and ascending links are copies of an arbitrary (finite) flag complex L (i.e. one with the property that if a collection of vertices is pairwise joined by edges then the collection bounds a simplex.[1] Note that the barycentric subdivision of any simplicial complex is a flag complex, so there is no restriction on the homeomorphism type. The example is built in the following way. Let G_L be the right-angled Artin group modelled on L, i.e. the group with a generator for each vertex of L, and the relations that two generators commute if the corresponding vertices are joined by an edge. Let $f : G_L \to \mathbb{Z}$ map each generator to 1, and let $H_L = Ker(f)$. When L is the n-fold join of a pair of points, H_L is the group from Example 7.3. What happens for different choices of L is discussed in [BB97]. The key is that there is a natural $K(G_L, 1)$, which is a torus complex (one torus for every simplex of L), and it can be used in the same way as the product of graphs was used in Example 7.2. Likewise, the infinite cyclic cover $K(H_L, 1)$ of $K(G_L, 1)$ admits a Morse function $K(H_L, 1) \to \mathbb{R}$ with descending and ascending links copies of L. For example, note that when L is contractible, the group H_L has type F since $K(H_L, 1)$ is homotopy equivalent to the quotient of a level set of the Morse function.

The following exercises may require looking up some definitions.

[1] Warren Dicks summarized the condition succinctly as "every non-simplex contains a non-edge".

Exercise 2. *Let T_n be the regular $(n+1)$-valent tree and let $h_n : T_n \to \mathbb{R}$ be a Busemann function corresponding to some end of T_n. The Diestel-Lieder graph $D(m, n)$ is defined as*

$$D(m, n) = \{(x, y) \in T_m \times T_n \mid h_m(x) + h_n(y) = 0\}$$

For example, $D(2, 2)$ is essentially the Cayley graph of the Lamplighter group

$$L = \mathbb{Z}_2 \wr \mathbb{Z} := \oplus_{i \in \mathbb{Z}} \mathbb{Z}_2 \ltimes \mathbb{Z}$$

where $\mathbb{Z}$ acts by shifting the summands (this is called the wreath product).

Prove that $D(m, n)$ is a connected graph. With a bit more work, show that if $m, n \geq 2$ then $D(m, n)$ has large loops, i.e. there are arbitrarily large loops that are not compositions of loops that tighten to loops shorter than the original. The Cayley graph of a group has this property if and only if the group is not finitely generated.

Exercise 3. *Prove that the asymptotic cone of the 3-dimensional solvable Lie group Sol is not simply connected. Hint: Sol is a horosphere in $\mathbb{H}^2 \times \mathbb{H}^2$, so its asymptotic cone is a horosphere in the product of two $\mathbb{R}$-trees.*

8. Finite subgroups

Consider the following question. Suppose G is a group that acts properly and cocompactly (and simplicially) on a contractible simplicial complex X. Are there only finitely many conjugacy classes of finite subgroups in G?

8.1. **Positive answers.**

It suffices to argue that every finite subgroup fixes a point – by simpliciality there must then be a fixed vertex, and by cocompactness there are only finitely many orbits of vertices, so every finite subgroup could be conjugated into one of finitely many vertex stabilizers.

A geometric situation where one can find fixed points is that of non-positive curvature. Suppose X is equipped with a $CAT(0)$ metric and G acts by isometries. Let $Q < G$ be a finite subgroup and choose a Q-orbit. There is a unique smallest closed metric ball that contains this orbit and its center is therefore fixed by Q.

Another situation where fixed points always exist is that of p-groups. Suppose $Q < G$ is a finite p-group for some prime p. Then by Smith theory Q fixes a point of X.

8.2. **Negative answers.** Until recently there were no examples known. The following construction is due to Leary-Nucinkis [LN03].

Let L be a finite flag complex and Q a finite group acting on L simplicially and without global fixed points. Let G_L and H_L be the groups from Example 7.4. Note that Q acts on both G_L and H_L so we may form $\tilde{G}_L := G_L \ltimes Q$ and $\tilde{H}_L := H_L \ltimes Q$. There is a homomorphism $\tilde{G}_L \to \mathbb{Z}$ whose kernel is H_L. The group $\tilde{G}_L$ acts on the universal cover X of the complex $K(G_L, 1)$ considered in Example 7.4. Each vertex of X is fixed by a copy of Q which is contained in $\tilde{H}_L$. There is only one $\tilde{G}_L$-orbit of vertices, so all these copies of Q are conjugate in $\tilde{G}_L$. However, the action of $\tilde{H}_L$ is height preserving, so two such copies of Q are conjugate if and only if the fixed vertices have the same height (by the condition that there are no global fixed points in L it follows that a copy of Q fixes a unique vertex). We conclude that there are infinitely many $\tilde{H}_L$-conjugacy classes of subgroups isomorphic to Q. In the orbifold language, one can see the argument as follows: The quotient $X/\tilde{H}_L$ is an "orbifold $K(\tilde{H}_L, 1)$" and there are infinitely many vertices whose label is a copy of Q. Since the edges have labels smaller than Q, it follows that the labels of different vertices are non-conjugate.

For example, we could take $Q = \mathbb{Z}_2$ and L the 3-fold join of a pair of points. Then the group $\tilde{H}_L$ is finitely presented and cointains infinitely many non-conjugate elements of order 2.

To get examples as in our Question, we need to take L to be contractible. For example, A_5 acts on a finite contractible 3-complex without global fixed points [FR59], leading to a counterexample with $Q = A_5$. Likewise, $Q = \mathbb{Z}_{30} \times \mathbb{Z}_{30}$ can be made to work as well [Oli75].

In fact, a bit more general construction (that allows L to be an infinite complex, but with a cocompact action of $\mathbb{Z}$) shows that there are examples of this sort whenever Q is not a p-group. See [Lea05].

References

[BB97] Mladen Bestvina and Noel Brady. Morse theory and finiteness properties of groups. *Invent. Math.*, 129(3):445–470, 1997.

[Bie76] Robert Bieri. *Homological dimension of discrete groups*. Mathematics Department, Queen Mary College, London, 1976. Queen Mary College Mathematics Notes.

[Bro94] Kenneth S. Brown. *Cohomology of groups*, volume 87 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994. Corrected reprint of the 1982 original.

[FR59] E. E. Floyd and R. W. Richardson. An action of a finite group on an n-cell without stationary points. *Bull. Amer. Math. Soc.*, 65:73–76, 1959.

[FS05] Daniel Farley and Lucas Sabalka. Discrete Morse theory and graph braid groups. *Algebr. Geom. Topol.*, 5:1075–1109 (electronic), 2005.

[Lea05] Ian J. Leary. On finite subgroups of groups of type VF. *Geom. Topol.*, 9:1953–1976 (electronic), 2005.

[LN03] Ian J. Leary and Brita E. A. Nucinkis. Some groups of type VF. *Invent. Math.*, 151(1):135–165, 2003.

[Mil63] J. Milnor. *Morse theory.* Based on lecture notes by M. Spivak and R. Wells. Annals of Mathematics Studies, No. 51. Princeton University Press, Princeton, N.J., 1963.

[Oli75] Robert Oliver. Fixed-point sets of group actions on finite acyclic complexes. *Comment. Math. Helv.*, 50:155–177, 1975.

[Sta63] John Stallings. A finitely presented group whose 3-dimensional integral homology is not finitely generated. *Amer. J. Math.*, 85:541–543, 1963.