Prologue: Trefoil Knot Group

The simplest non-trivial knot is the trefoil knot shown in Figure 1. As a way to introduce the flavor of geometric group theory we ask: *What can we say about the fundamental group of its complement?* Our primary goal is to illustrate how geometric arguments can be used to prove purely algebraic results. The arguments are merely sketched, but the reader should be able to go back and fill in the details as they work their way through the text.

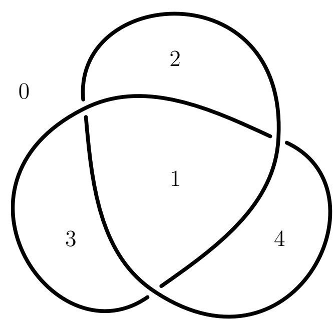

FIGURE 1. The trefoil knot.

To establish notation, let K denote the knot shown and let G be the group $\pi_1(\mathbb{S}^3 \setminus K)$. (For technical reasons it is cleaner and more symmetric to work in $\mathbb{S}^3$, the 1-point compactification of $\mathbb{R}^3$, than in $\mathbb{R}^3$ itself.) The first thing to notice is that *G is the fundamental group of a compact two-dimensional complex.* To show this we construct a 2-complex $\mathcal{D}$ inside $\mathbb{S}^3 \setminus K$ and then deform $\mathbb{S}^3 \setminus K$ down to $\mathcal{D}$. Since deformation retractions do not alter fundamental groups, $G = \pi_1(\mathbb{S}^3 \setminus K) \approx \pi_1(\mathcal{D})$. There are two common constructions for retracting arbitrary knot complements onto two-dimensional subspaces, usually attributed to Dehn and Wirtinger. Since we are using Dehn's procedure, the final result is known as a *Dehn complex*.

The Dehn complex for the trefoil knot shown in Figure 1 has two vertices, five edges and three 2-cells. To construct it we think of K as living in a small neighborhood of the xy-plane (or rather in a small neighborhood of its 1-point compactification, an equatorial 2-sphere inside $\mathbb{S}^3$), and we place a vertex v_+ above K and a vertex v_- below K. Next, we add an edge for each of the five regions of the xy-plane determined by the projection of K. More concretely, the regions in Figure 1 have been numbered and we add an edge E_i that connects v_- and v_+, oriented from v_- to v_+, passing through region i. Finally, we add a 2-cell for each of the three crossings. Each 2-cell is a square folded to look like Figure 2. The boundary of the square is then identified with the four edges corresponding to the four adjacent regions. The crossing at the top left of Figure 1, for example, creates

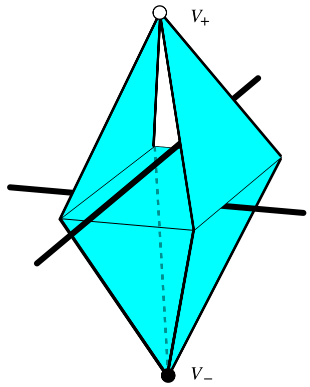

FIGURE 2. The 2-cell at a crossing

a square whose boundary follows the path $E_0^{-1}E_3E_1^{-1}E_2$. This can be interpretated as claiming that the loop $E_0^{-1}E_3$ based at v_+ is homotopic to the loop $E_2^{-1}E_1$. The other two 2-cells are attached along $E_0^{-1}E_4E_1^{-1}E_3$ and $E_0^{-1}E_2E_1^{-1}E_4$, respectively. See Figure 3.

The deformation retraction from $\mathbb{S}^3 \setminus K$ to $\mathcal{D}$ alluded to above expands away from K like adding air into a long balloon. Parts of this retraction are easy to visualize. In Figure 2, for example, the complement of K inside this tent clearly retracts onto the folded square. Piecing together these local pictures, we find that G is the fundamental group of $\mathcal{D}$ and, as a consequence, that G acts freely on its universal cover $\widetilde{\mathcal{D}}$.

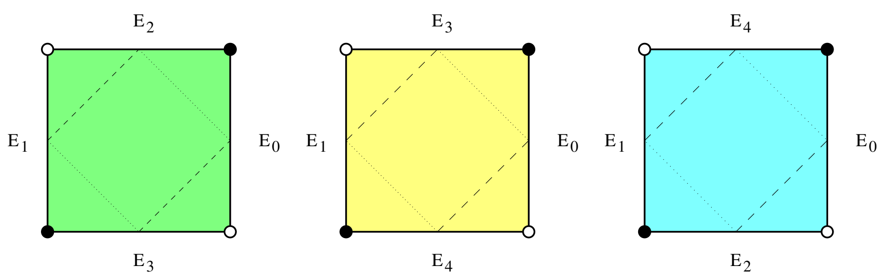

FIGURE 3. The three 2-cells in $\mathcal{D}$. The open circles represent v_+ and the closed circles represent v_-. To make these 2-cells look like the one shown in Figure 2, fold up along the dashed lines and down along the dotted ones.

The next key idea is that *if we understand the geometry of $\widetilde{\mathcal{D}}$ and the way G acts on it, then we gain insight into the the algebraic structure of G as a group.* The geometry of $\widetilde{\mathcal{D}}$ is quite elegant. Since $\mathcal{D}$ contains only three 2-cells, $\widetilde{\mathcal{D}}$ has only three equivalence classes of 2-cells under the action of G. For convenience we refer to these as the green, yellow and blue 2-cells, reading left to right in Figure 3. The edges E_0 and E_1 are both contained in all three 2-cells, while the other three edges only occur in two of the three 2-cells. In fact, if you fix a particular lift of the green 2-cell in $\widetilde{\mathcal{D}}$, oriented as shown, then there is a unique yellow 2-cell below it, followed by a unique blue 2-cell, followed by a unique green 2-cell, and so on. The two sides of this infinite strip consist of lifts of the edges E_0 and E_1, alternating on both sides. With a bit more work one sees that the local structure of $\widetilde{\mathcal{D}}$ looks like Figure 4 and that as a topological space $\widetilde{\mathcal{D}}$ is homeomorphic to the direct product of the real line and an infinite, trivalent tree.

Actually, even more is true. We can add a metric to $\widetilde{\mathcal{D}}$ by making each 2-cell isometric to a unit Euclidean square. The metric space $\widetilde{\mathcal{D}}$ still splits as a direct product, this time of the real line with the standard metric and a metric trivalent tree where each edge has length 1. In other words, if we let $\mathcal{T}_3$ denote the infinite, trivalent tree with edges of unit length, then $\widetilde{\mathcal{D}}$ is isometric to $\mathcal{T}_3 \times \mathbb{R}$. The action of G preserves the metric as well as the product structure on $\mathcal{T}_3 \times \mathbb{R}$ so that by projecting onto the first or the second factor, the group G acts by isometries on $\mathcal{T}_3$ and it acts by isometries on $\mathbb{R}$.

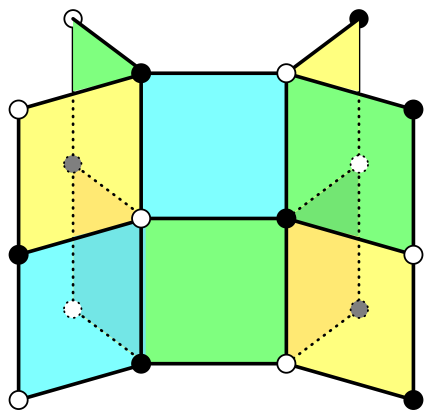

FIGURE 4. The local structure in $\widetilde{\mathcal{D}}$.

The last bit of preparation we need is to find a presentation of the group G. Any $g \in G$, acting on $\widetilde{\mathcal{D}}$, takes lifts of v_+ to lifts of v_+. In order to get a generating set for G it suffices to pick enough elements of G so that any lift of v_+ can be moved to any other using some composition of the actions of these elements and their inverses. Let v be a fixed lift of v_+ in $\widetilde{\mathcal{D}}$, and let a, b and c represent the unique elements of G that move v diagonally up and across the unique green, yellow and blue squares, respectively, that have v as a bottom corner. To see that a, b and c generate G, let V be the orbit of v under the action of the subgroup generated by a, b and c. Suppose that u is in V, g is the element of G that sends v to u, and u' is a lift of v_+ connected to u along the diagonal of a single square. Since g sends the vertices connected to v by a diagonal to the vertices connected to u by a diagonal, we can find an element h in the set $\{a, b, c, a^{-1}, b^{-1}, c^{-1}\}$ so that gh sends v to u'. Notice that we are *precomposing* g with h which involves *right* multiplication by h. This is because we always assume that our groups act on the left. See Appendix A. In any case, this shows that the vertices in V are closed under adjacency. Geometrically, it is now clear that every lift of v_+ lies in V, and thus a, b and c generate all of G.

Finally, suppose that v is the open circle on the bottom of Figure 4 slightly to the right of center. The reader can check that the words ab, bc and ca all move v to the open circle directly above it, so these words all represent the same element in G. When projected to $\mathcal{D}$, the products ab, bc and ca are homotopic, and represent the loop that passes through the central region of the trefoil knot and then returns to v_+ via the exterior region. With only a bit more work, one can show that the presentation $\langle a, b, c \mid ab = bc = ca\rangle$ is a presentation for G.

The group G acts freely and cocompactly on a contractible complex $\widetilde{\mathcal{D}}$. We understand the structure of $\widetilde{\mathcal{D}}$ and the action of G, and we have "words" we can use to describe the elements of G. We can now establish the following:

THEOREM. *If G is the fundamental group of the trefoil knot complement then*

1. *We can efficiently determine whether a word in the generators represents the identity;*
2. *The group G contains no nontrivial element of finite order;*
3. *The kernel of the map $f : G \twoheadrightarrow \mathbb{Z}$ sending a, b and c to $1 \in \mathbb{Z}$ is a free group of rank two;*
4. *The element $z = (ab)^3 = (bc)^3 = (ca)^3$ is central in G;*
5. *The group G contains a finite index subgroup isomorphic to $\mathbb{F}_2 \times \mathbb{Z}$;*
6. *The group G is residually finite, meaning that the intersection of all finite index subgroups of G is the trivial subgroup $\{1\}$.*
7. *The element z generates the center, so that $Z(G) = \langle z\rangle$;*
8. *The quotient $G/Z(G)$ is isomorphic to $PSL_2(\mathbb{Z})$;*

How are such claims established? We outline one approach and completely ignore the technical details.

SKETCH OF PROOF: To tell whether an element in the generators represents the identity, simply trace out its effect on a lift of v_+ inside $\widetilde{\mathcal{D}}$. If this lift ends where it started then this word represents the identity; otherwise, it does not. The reason this works is because the process of constructing the universal cover $\widetilde{\mathcal{D}}$ secretly encodes a solution to the "word problem" for G. See Chapter 3 for details.

To prove 2 we combine the action of G on the factors of $\widetilde{\mathcal{D}} \approx \mathcal{T}_3 \times \mathbb{R}$ with the fact that any finite order isometry of a metric tree must fix a point. (This fact is proved in Chapter 5.) Thus, any $g \in G$ of finite order fixes a point in $\mathcal{T}_3$ and it fixes a point in $\mathbb{R}$, so it fixes a point in $\widetilde{\mathcal{D}}$. But the action of G on $\widetilde{\mathcal{D}}$ is free so g is the identity.

Let H be the kernel of the map $f : G \to \mathbb{Z}$ described in item 3. Since the action of H on $\widetilde{\mathcal{D}}$ projects to a free action on the tree $\mathcal{T}_3$, the fundamental group of the quotient of $\mathcal{T}_3$ by this action is isomorphic to H. This quotient has two vertices, three edges and its fundamental group is $\mathbb{F}_2$.

The action of the element $z = (ab)^3 = (bc)^3 = (ca)^3$ on $\widetilde{\mathcal{D}}$ is a rotationless vertical translation. It can then be checked that pre- and post-composing any $g \in G$ with z results in the same action on $\widetilde{\mathcal{D}}$, hence both expressions describe the same element of G. (Actually, it is sufficient to check that this is true for a, b, and c since they generate G.) This proves 4. Item 5 is now immediate since the subgroup generated by H and z is isomorphic to $\mathbb{F}_2 \times \mathbb{Z}$ and index 6 in G.

Next, the easiest way to prove 6 is to combine item 5 with two easily proved facts: free groups are residually finite, and the class of residually finite groups is closed under direct product and finite extension.

To prove 7, let v be a particular lift of v_+ inside $\widetilde{\mathcal{D}}$ and consider the orbit of v under the action of $Z(G)$. Because of the symmetry of the situation with respect to a, b and c, the orbit of $Z(G)$ must be invariant under a $2\pi/3$ rotation around the vertical line through v. On the other hand, since the free group of rank 2 has trivial center, $Z(G) \cap H$ only contains the identity element and thus the orbit of v can have at most one element at each height. Combining these two ideas shows that 1) the orbit of v under the action of $Z(G)$ is contained in the vertical line through v, and 2) $Z(G)$ must be generated by the element that produces the smallest possible positive vertical change when applied to v. By 4, z is central and it moves v up six steps. There are exactly two elements of G, namely ab and $(ab)^2$, that move v to a lift of v_+ that is both on the vertical line through v and between v and its image under z. After checking that ab and $(ab)^2$ are not central, we conclude that $Z(G) = \langle z \rangle$.

Finally, to prove 8 we note that the action of $Z(G)$ on the $\mathcal{T}_3$ factor is trivial. Thus, we get a well-defined action of the quotient group $G/Z(G)$ on the trivalent tree $\mathcal{T}_3$. There is a well-known action of $\mathrm{PSL}_2(\mathbb{Z})$ on $\mathcal{T}_3$, and, by comparing the two actions, we can see that the groups are identical. □

Exercises

1. (Details) Fill in as many of the details of the proof of the theorem as you can. Alternatively, make a list of the arguments that seem unclear to you or imprecise at this point.
2. (Figure 8 knot) Let K be the knot shown in Figure 5.
   a. Construct the Dehn complex $\mathcal{D}$ for K.
   b. Draw a small portion of $\widetilde{\mathcal{D}}$ and try to understand its structure. Be forewarned that this is more difficult than it was for the trefoil knot.

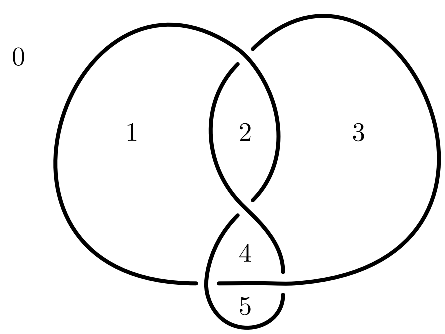

FIGURE 5. The figure 8 knot