The Homology of Richard Thompson's Group F

Kenneth S. Brown

Dedicated to Ross Geoghegan in honor of his 60th birthday.

Introduction

Let F be Richard Thompson's group, which can be defined by the presentation

$$F = \langle\, x_0, x_1, x_2, \ldots \,;\, x_n^{x_i} = x_{n+1} \text{ for } i < n \,\rangle \ .$$

Here $x^y := y^{-1}xy$. One can also describe F in a variety of other ways, some of which are reviewed briefly in Section 1. In the early 1980s Ross Geoghegan and I studied the homological properties of F; see [6–8]. We showed, among other things, that the integral homology $H_*(F)$ is free abelian of rank 2 in every positive dimension. It turns out that the homology admits a natural ring structure, which I calculated a few years later. The answer is that $H_*(F)$ is an associative ring (without identity) generated by an element ε of degree 0 and elements α, β of degree 1, subject to the relations

$$\varepsilon^2 = \varepsilon$$
$$\varepsilon\alpha = \beta\varepsilon = 0$$
$$\alpha\varepsilon = \alpha\ , \quad \varepsilon\beta = \beta\ .$$

It follows that $\alpha^2 = \beta^2 = 0$ and that the alternating products $\alpha\beta\alpha\cdots$ and $\beta\alpha\beta\cdots$ give a basis for the homology in positive dimensions.

With the aid of this ring structure on the homology, one can calculate the integral cohomology ring:

$$(0.1) \qquad H^*(F) \cong \bigwedge(a,b) \otimes \Gamma(u)\ ,$$

where $\bigwedge(a,b)$ is an exterior algebra on two generators a, b of degree 1, and $\Gamma(u)$ is a divided polynomial ring on one generator u of degree 2.

I never published these results because they were subsumed by the work of Melanie Stein [20], who proved analogous results for a much larger class of groups. Since F remains of great current interest, and since readers may find it inconvenient to work through Stein's paper and specialize everything to the case of F, I give

here my original proofs for that case. In particular, I explain the cohomology calculation (0.1), which is not stated explicitly in [20].

The impetus for publishing these 15-year-old results at the present time came from a question recently raised by Geoghegan [private communication]: Is F a "Kähler group", i.e., the fundamental group of a compact Kähler manifold? The hope was that knowledge of the ring structure on $H^*(F)$ would yield a negative answer to Geoghegan's question. This idea did not work, but Napier and Ramachandran [16] have since used completely different ideas to show that, in fact, F is not a Kähler group.

The remainder of the paper is organized as follows. Section 1 is a brief review of F and some of its properties. I introduce in Section 2 a homomorphism $\mu\colon F \times F \to F$, which induces the product structure on $H_*(F)$ alluded to above. The product μ is not strictly associative, but it is associative up to conjugacy, so the homology becomes an associative ring. In Section 3 I describe a cubical $K(F,1)$-complex X due to Stein. It is a variant of a $K(F,1)$ constructed in [7], but it has the advantage that it comes equipped with a strictly associative product $X \times X \to X$, which induces the product μ on π_1. [Note: One has to be careful about basepoints in order to make sense of the last statement, since the product is not basepoint-preserving.] The cellular chain complex $C_*(X)$ then becomes a differential graded ring. Its structure is so simple that one can compute its homology directly; this is done in Section 4. Finally, I explain in Section 5 how the homology calculation yields the cohomology result stated in (0.1).

This paper is an expanded version of a talk I gave at an AMS special session in Nashville on October 17, 2004, in honor of Ross Geoghegan's 60th birthday. Ross is the person who first told me about Thompson's group and made me realize what a fascinating object it is. It is a great pleasure to dedicate this paper to him.

1. Background on F

There are many references for the basic facts about F, including [1–3, 7, 9], all of which have further references in their bibliographies. We review in this section only a few basic facts that we will need later.

1.1. Dyadic PL-homeomorphisms. Let I and J be intervals of real numbers. A homeomorphism $f\colon I \to J$ will be called *dyadic* if it satisfies the following conditions:

- f is piecewise linear with only finitely many breakpoints.
- Every breakpoint has dyadic rational coordinates.
- Every slope is an integral power of 2.

A basic fact about F, essentially known to Thompson, is that F can be identified with the group of dyadic homeomorphisms of the unit interval $[0, 1]$, with the group law being composition: $(fg)(t) = f(g(t))$ for $f, g \in F$ and $t \in [0, 1]$. [Warning: Some authors, including Thompson himself, have used the opposite convention for composition.] Figure 1 shows the graphs of the first two generators, x_0 and x_1. Note that x_1 is the identity on $[0, 1/2]$ and is a rescaled copy of x_0 on $[1/2, 1]$. The remaining generators $x_2, x_3, \ldots$ are obtained by repeating this process. More precisely, there is a homomorphism $\phi\colon F \to F$ such that $\phi(f)$ is the identity on $[0, 1/2]$ and is a rescaled copy of f on $[1/2, 1]$, and the assertion is that $\phi(x_n) = x_{n+1}$ for all $n \geq 0$.
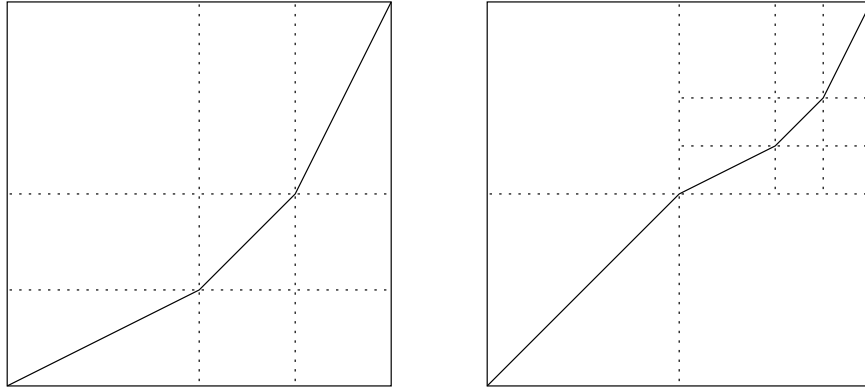
FIGURE 1. The generators x_0 and x_1

All dyadic PL-homeomorphisms of $[0, 1]$, and hence all elements of F, can be obtained from *dyadic subdivisions*, as illustrated by the dotted lines in Figure 1. One subdivides the domain and range into an equal number of parts by repeated insertion of midpoints, and one maps the subintervals of the domain subdivision linearly to the subintervals of the range subdivision.

Finally, we recall for future reference that there is a fairly obvious way to represent dyadic subdivisions by rooted binary trees, as explained in [1–3, 20]. More generally, we will have occasion to talk about dyadic subdivisions of $[0, n]$, where n is an integer ≥ 1, in which we start with the subdivision into unit intervals $[i - 1, i]$ $(i = 1, \ldots, n)$ and then dyadically subdivides these, as in Figure 2. Such a

FIGURE 2. Two dyadic subdivisons of $[0, 2]$

subdivision is represented by a binary forest consisting of n rooted binary trees (in a definite order). Thus the roots of the forest correspond to the intervals $[i - 1, i]$, and the leaves correspond to the parts of the sudivision. For example, the two subdivisions of $[0, 2]$ in Figure 2 are represented by the forests in Figure 3.
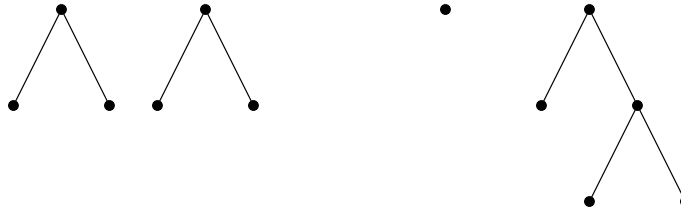

FIGURE 3. Forests represent dyadic subdivisions

1.2. Conjugacy idempotents. Dydak [11, 12] and, independently, Freyd and Heller [13] rediscovered Thompson's group in connection with a problem in homotopy theory. Briefly, they were interested in "free homotopy idempotents", i.e.,

maps from a space to itself that are idempotent up to homotopy; the maps are required to be basepoint preserving, but the homotopy is not. Passing to the fundamental group, one is led to study group homomorphisms that are idempotent up to conjugacy, and it turns out that F carries the universal example of such a homomorphism. Namely, the defining relations for F show that the shift homomorphism ϕ discussed above satisfies

$$\phi^2(f) = \phi(f)^{x_0}$$

for all $f \in F$, so ϕ is idempotent up to conjugacy. To see universality, suppose we have a group G and an endomorphism $\psi \colon G \to G$ such that $\psi^2 = \psi^{y_0}$ for some $y_0 \in G$. Then G contains elements $y_n := \psi^n(y_0)$ for $n \geq 1$. The equation $\psi^2 = \psi^{y_0}$ implies that $y_j^{y_i} = y_{j+1}$ for $j > i = 0$, and we can repeatedly apply ψ to see that this remains valid for $j > i \geq 0$. Thus there is a homomorphism $F \to G$ taking x_n to y_n for all n.

1.3. Algebra automorphisms. Consider the algebraic system consisting of a set A together with a bijection $A \times A \to A$. Thus there is a product on A, and every $a \in A$ factors uniquely as a product of two other elements. Algebras of this type have appeared in several places; see [10, 14, 15, 19, 21]. Following Smirnov [19], we call A a *Cantor algebra*. Galvin and Thompson [unpublished] showed that F is isomorphic to the group of "order-preserving" automorphisms of the free Cantor algebra on one generator. This point of view was exploited in [3]. In order to explain what "order-preserving" means, we need to recall some facts about bases of free Cantor algebras. The definitions will be adapted to our present needs, in which bases are always ordered.

All Cantor algebras considered in this paper will come equipped with a "standard basis" $a_1, \ldots, a_n$, in which the order of the basis elements is important. A *simple expansion* of an ordered basis consists of factoring one of the basis elements a as $a_0 a_1$ and replacing a by the two elements a_0, a_1 (in that order). A general expansion consists of doing finitely many simple expansions. It is an easy fact, proved in several of the references cited above, that an expansion of a basis is again a basis. For example, if A is the free Cantor algebra on one generator a, then we can factor a as a triple product in two different ways:

$$a = a_0 a_1 = a_0(a_{10} a_{11}) , \qquad a = a_0 a_1 = (a_{00} a_{01}) a_1 .$$

This yields a basis a_0, a_1 of size 2 and two bases a_0, a_{10}, a_{11} and a_{00}, a_{01}, a_1 of size 3. The opposite of expansion is called *contraction*. A *simple contraction* consists of replacing two consecutive basis elements by their product, and a general contraction consists of finitely many simple contractions. Finally, an *ordered basis* of our Cantor algebra is one that can be obtained from the standard basis by doing finitely many expansions or contractions.

REMARK 1.1. Expanding a basis is analogous to subdividing an interval. In particular, if we start with an ordered basis having n elements, then there is a fairly obvious way to represent a k-fold expansion of it by a binary forest with n roots and $n + k$ leaves. For example, given an ordered basis a, b, we have a 2-fold expansion a_0, a_1, b_0, b_1 and a second 2-fold expansion a, b_0, b_{10}, b_{11}. These are represented by the two forests in Figure 3 above.

Given two Cantor algebras of the type we are considering (free with a given linearly ordered basis), we define an *order-preserving* isomorphism between them

to be one that takes an ordered basis of the domain to an ordered basis of the range (preserving the ordering on the bases). For example, if A is free on a as above, then there is an automorphism of A taking a_0, a_{10}, a_{11} to a_{00}, a_{01}, a_1. This corresponds to x_0. Using the analogy between subdivision and expansion, the reader should be able to look at Figure 1 and guess which automorphism of A corresponds to x_1.

The ordered bases of the free Cantor algebra A on one generator form a poset $\mathcal{B}$, in which $B \leq C$ if C is an expansion of B. This poset is a directed set; it played an important role in [3] and will be referred to again later.

REMARK 1.2. For readers who prefer to avoid Cantor algebras, here is an alternate description of the poset $\mathcal{B}$ of ordered bases. Consider dyadic PL-homeomorphisms $f\colon [0, n] \to [0, 1]$ where n is an integer ≥ 1. These play the role of bases. [To see why, note that giving an ordered basis of A of size n is equivalent to giving an order-preserving isomorphism $A_n \to A$, where A_n is free on $a_1, \ldots, a_n$.] An *expansion* of f is a homeomorphism $[0, n + k] \to [0, 1]$ of the form $f \circ s$, where $s\colon [0, n + k] \to [0, n]$ is a *subdivision map*, obtained as follows: Perform a dyadic subdivision of $[0, n]$ into $n + k$ parts, and let s map the unit intervals $[j - 1, j]$ of $[0, n + k]$ to the $n + k$ subintervals of $[0, n]$. See Figure 4 for an example with $n = 1$ and $k = 2$.
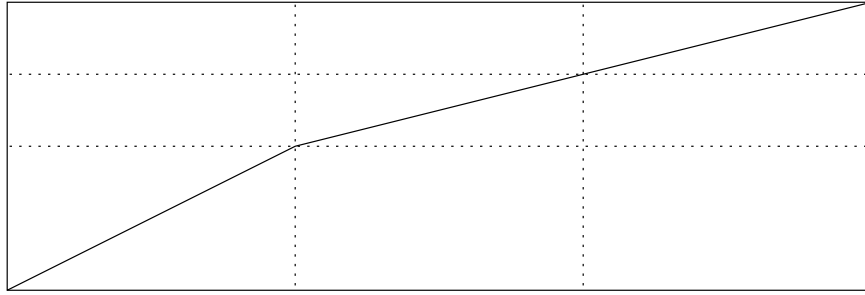

FIGURE 4. A subdivision map $[0, 3] \to [0, 1]$

1.4. Finiteness properties of F. It is obvious that F is generated by the first two generators x_0, x_1, since the remaining x_n are obtained by repeated conjugation. Less obviously, two relations suffice:

$$x_1^{x_0 x_0} = x_1^{x_0 x_1}$$
$$x_1^{x_0 x_0 x_0} = x_1^{x_0 x_0 x_1} \ .$$

This pattern was extended in Brown–Geoghegan [7], where it was shown that there is an Eilenberg–MacLane complex Y of type $K(F, 1)$ with exactly two cells in each positive dimension. Our method was motivated by the connection between F and the theory of free homotopy idempotents. Namely, we constructed the universal example of a space X with a free homotopy idempotent, and we showed that X was a $K(F, 1)$-complex. Now X has infinitely many cells in each positive dimension, but we were able to show, by imitating the proof that F requires only two generators and two relations, that there were only two "homotopically essential" cells in each dimension. The desired complex Y was then obtained as a quotient complex of X.

Incidentally, this is where the homology calculation cited in the introduction came from: We showed that the cellular chain complex $C_*(Y)$ has trivial boundary operator, so that $H_n(F) \cong \mathbb{Z}^2$ for each $n \geq 1$.

1.5. The abelianization of F. Let F' be the commutator subgroup of F and let F_{ab} be the abelianization F/F'; it is isomorphic to $\mathbb{Z} \times \mathbb{Z}$. From the point of view of dyadic PL-homeomorphisms, the abelianization map $F \to \mathbb{Z} \times \mathbb{Z}$ is given by $f \mapsto \big(\log_2 f'(0), \log_2 f'(1) \big)$. [Recall that the slopes of elements of F are in the infinite cyclic multiplicative group generated by 2.] Thus F' consists of the elements of F that are the identity near the endpoints $0, 1$. The group F' is simple but infinitely generated; it is the union of an increasing sequence of isomorphic copies of F.

2. A product on F

2.1. Definition. There is a homomorphism $\mu \colon F \times F \to F$, denoted $(f, g) \mapsto f * g$, which is defined as follows: If we interpret F as the group of dyadic homeomorphisms of $[0, 1]$, then $f * g$ is given by

$$(f * g)(t) = \begin{cases} f(2t)/2 & 0 \leq t \leq 1/2 \\ (g(2t - 1) + 1)/2 & 1/2 \leq t \leq 1. \end{cases}$$

Less formally, $f * g$ is a rescaled copy of f on $[0, 1/2]$ and a rescaled copy of g on $[1/2, 1]$; see Figure 5 for an example. Alternatively, if we interpret F as the group
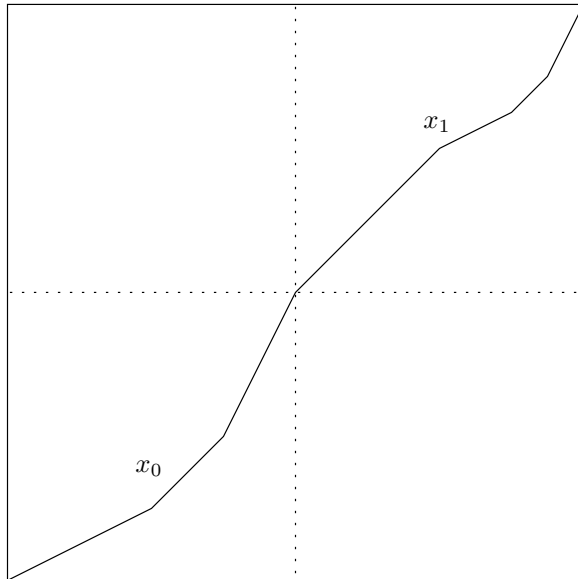

FIGURE 5. $x_0 * x_1$

of order-preserving isomorphisms of the free Cantor algebra A on one generator a, then we can first form the automorphism $f \amalg g$ of $A \amalg A$, where $A \amalg A$ is the categorical sum of two copies of A and hence is free on two generators, and then we can transport this to an automorphism of A via a suitable isomorphism $A \cong A \amalg A$.

One can check that μ is associative up to conjugacy:

$$(2.1) \qquad\qquad f * (g * h) = ((f * g) * h)^{x_0} .$$

There is, however, no identity. In particular, the identity element 1 for the group law on F is not an identity for μ; indeed, one has

$$(2.2) \qquad\qquad 1 * f = \phi(f) ,$$

where ϕ is the conjugacy idempotent discussed above. Similarly, μ-multiplying on the right by 1 is a (different) conjugacy idempotent on F.

2.2. The induced product on homology. Exactly as in the homology theory of abelian groups [5, V.5], the product μ induces a product

$$H_*(F) \otimes H_*(F) \to H_*(F) ,$$

making $H_*(F)$ an associative ring, where associativity comes from (2.1) and the fact that inner automorphisms act trivially on homology [5, II.6.2]. This product has no identity; the canonical generator ε of $H_0(F) = \mathbb{Z}$ is idempotent, and left multiplication by it is the endomorphism ϕ_* of $H_*(F)$ by (2.2). As we will see, this idempotent endomorphism has rank 1 in every positive dimension, as does the encomorphism given by right multiplication by ε.

We now wish to calculate the product explicitly. We will do this in Section 4 after recalling a construction due to Stein [20].

3. Stein's cubical $K(F,1)$

3.1. A groupoid analogue of F. Let $\mathcal{F}$ be the category whose objects are the intervals $[0, n]$ ($n \geq 1$) and whose morphisms are the dyadic PL-homeomorphisms between them. Alternatively, we could take the objects to be the free Cantor algebras A_n and the morphisms to be the order-preserving isomorphisms. Let $|\mathcal{F}|$ denote the geometric realization of $\mathcal{F}$ as defined by Quillen [17]. See also [18], where the theory is reviewed for the case where the category is a poset. [A poset can be viewed as a category, with one morphism $a \to b$ for every relation $a \leq b$.] Recall that $|\mathcal{F}|$ is a CW-complex in which every open p-cell can be identified with the interior of the standard p-simplex. There is one such p-cell for every p-tuple of composable non-identity morphisms in $\mathcal{F}$:

$$I_{n_0} \xrightarrow{f_1} I_{n_1} \xrightarrow{f_2} \cdots \xrightarrow{f_p} I_{n_p} .$$

Faces are gotten by deleting objects and, if the object is not the first or last, composing morphisms. [It is possible that this will yield an identity map, hence a "degenerate" simplex that must be collapsed to a lower-dimensional simplex.] Since $\mathcal{F}$ is a connected groupoid and the group of automorphisms of one object is F, it follows from Quillen's theory that $|\mathcal{F}|$ is an Eilenberg–MacLane complex of type $K(F,1)$. We will not actually need to make use of this fact, but it provides helpful motivation.

3.2. A smaller category. The complex $\mathcal{F}$ is much too big to be of any use computationally, so we pass to a subcategory $\mathcal{S}$ such that $|\mathcal{S}|$ is still a $K(F,1)$. The objects of $\mathcal{S}$ are the same as those of $\mathcal{F}$, but as morphisms we only use the subdivision maps s as defined in Remark 1.2 in Section 1.3. Alternatively, if we want to use Cantor algebras as the objects, we only use isomorphisms $A_{n+k} \to A_n$ that take the standard basis elements of A_{n+k} to basis elements of A_n obtained by

doing a k-fold expansion of its standard basis. From either point of view, one sees immediately that the morphisms in $\mathcal{S}$ from the object associated with $n + k$ to the object associated with n correspond to binary forests with n roots and $n+k$ leaves. See Belk [1, Section 7.2] for further remarks on $\mathcal{S}$.

Although we will not need this fact, one can check that $|\mathcal{S}|$ is a regular CW-complex in which every closed cell is canonically homeomorphic to a standard simplex. It is not a simplicial complex, however, because a simplex is not determined by its vertices. But we will see shortly that its universal cover is a simplicial complex.

To see that $|\mathcal{S}|$ is a $K(F, 1)$-complex, we can proceed in two different ways. The first method, based on [17, Theorem A], is to consider the fibers in the sense of Quillen of the inclusion $\mathcal{S} \hookrightarrow \mathcal{F}$; it suffices to show that they are contractible. A direct check of the definitions shows that each fiber is a poset and is a directed set, hence it is indeed contractible. For example, the fiber over I_1 is isomorphic to the poset $\mathcal{B}$ discussed in Section 1.3, with the order relation reversed.

A more elementary approach is to directly construct the universal cover of $|\mathcal{S}|$ and observe that it is contractible. In fact, the universal cover turns out to be the geometric realization $|\mathcal{B}|$ of the contractible poset $\mathcal{B}$. One can see this by thinking of $\mathcal{B}$ either as the poset of bases or as a poset constructed using dydadic maps of intervals. From either point of view there is an obvious action of F on $\mathcal{B}$, and it is straightforward to check that the induced action on the (contractible) simplicial complex $|\mathcal{B}|$ is free and that $|\mathcal{S}|$ is the quotient. Hence $|\mathcal{S}|$ is indeed a $K(F, 1)$-complex.

3.3. A cubical complex. Following Stein [20], we now pass to a further subcomplex $X \subset |\mathcal{S}|$, which is still a $K(F, 1)$ and in which the simplices can be naturally grouped into cubes. It is easiest to carry this out in the universal cover $|\mathcal{B}|$ and then pass to the quotient by the action of F. We will work with the original definition of $\mathcal{B}$ as the poset of ordered bases of $A = A_1$. Readers who prefer to work with intervals and dyadic subdivisions can translate everything into that language or can refer to [20].

Given an ordered basis L, an *elementary expansion* of L is a basis M obtained by replacing zero or more elements $b \in L$ by their factors b_0, b_1. Such an expansion corresponds to an "elementary forest", by which we mean one in which each tree consists either of the root only or the root with a single pair of descendants. For example, the expansion represented by the forest on the left in Figure 3 is elementary, while the one on the right is not. We will write $L \preceq M$ if M is an elementary expansion of L. [Warning: This relation is not transitive.] Recall that a simplex of $|\mathcal{B}|$ is given by a chain $L_0 < L_1 < \cdots < L_p$ of bases. Call the simplex *elementary* if $L_0 \preceq L_p$. This implies that $L_i \preceq L_j$ for $0 \le i \le j \le p$. Hence every face of an elementary simplex is elementary, and the elementary simplices form an F-invariant subcomplex $\tilde{X}$ of $|\mathcal{B}|$. Passing to the quotient by the action of F, we obtain the desired subcomplex $X \subset |\mathcal{S}|$. It has one cell for each chain of bases $L_0 < \cdots < L_p$ such that L_0 is the standard basis of A_n for some $n \ge 1$ and L_p is an elementary expansion of L_0.

Stein proves that the complex of elementary simplices is contractible, so that X is again a $K(F, 1)$-complex. Her proof is given in detail in [20] and is repeated, in a slightly different context, in [4], so we will not repeat it again here.

Finally, we will explain, still following Stein, how to give a coarser cell decomposition of X by lumping certain simplices into cubes, one for each elementary

expansion $L \preceq M$ in which L is the standard basis of some some A_n. Once again, it is easier to first do this in the universal cover, so we consider $L \preceq M$ with arbitrary $L \in \mathcal{B}$ and with M a k-fold elementary expansion of L. Then the interval $[L, M]$ in $\mathcal{B}$ is isomorphic, as a poset, to $\{0, 1\}^k$, where $0 < 1$ and the product is ordered component-wise. Since the geometric realization of $\{0, 1\}$ is canonically homeomorphic to the unit interval $[0, 1]$, we conclude that the geometric realization $\big|[L, M]\big|$ is a simplicially subdivided k-cube.

The relative interior of this k-cube is the union of the open simplices corresponding to the chains $L_0 < \cdots < L_p$ with $L_0 = L$ and $L_p = M$, so these interiors partition $\tilde{X}$. We therefore have a decomposition of $\tilde{X}$ as a regular cell complex in which all the closed cells are cubes. Passing to the quotient by F, we obtain the desired decomposition of X. The closed cells are cubes with some identifications on the boundary. For example, there is a 2-cube in X for every $1 \le i < j \le n$, corresponding to the following diagram in $\mathcal{S}$:

$$
\begin{array}{ccc}
A_{n+2} & \xrightarrow{\ s_i\ } & A_{n+1} \\
\ \downarrow{\scriptstyle s_{j+1}} & \searrow & \ \downarrow{\scriptstyle s_j} \\
A_{n+1} & \xrightarrow[\ s_i\ ]{} & A_n
\end{array}
$$

Here the top s_i is the isomorphism that maps the standard basis of A_{n+2} to the simple expansion at position i of the standard basis of A_{n+1}; the other maps are defined similarly. The diagram reflects the fact that there are two ways to expand the standard basis of A_n at positions i, j: First expand at position j and then expand the ith basis element of the result, or first expand at position i and then expand the $(j + 1)$st basis element of the result.

3.4. A product on X. There is a product $\mu \colon \mathcal{F} \times \mathcal{F} \to \mathcal{F}$, analogous to the product on F defined in Section 2.1, except that no rescaling is required. From the point of view of dyadic PL-homeomorphisms, we define μ on objects by setting $I_n * I_m = I_{n+m}$, and we define it on maps by gluing them together in the obvious way. Thus if f has domain I_n and g has domain I_m, then $f * g$ is f on $[0, n]$ and a shifted copy of g on $[n, n + m]$. Alternatively, in terms of Cantor algebras, we set $A_n * A_m = A_n \amalg A_m$, which we identify with A_{n+m}, and we set $f * g = f \amalg g$.

Notice that the product μ is strictly associative; this is the advantage of working with the groupoid $\mathcal{F}$ instead of the group F. Thus $|\mathcal{F}|$ becomes a topological semigroup, with $|\mathcal{S}|$ and X as subsemigroups. If we use the cubical structure on X, we find that the product of two cubes is again a cube, so that the set of cubes is a semigroup. This semigroup is quite easy to understand: Recalling that cubes correspond to elementary expansions, which are represented by elementary forests, the product is given by taking the disjoint union of the forests. It follows at once that the semigroup of cubes is the free semigroup on two generators v, e, where v is the vertex corresponding to the object A_1 and e is the unique edge in X from v to v^2. [Note that the 0-cube v is represented by the trivial forest with one root, and the 1-cube e is represented by the forest consisting of a root that has two descendants.] For example, the 2-cube corresponding to the elementary expansion shown in Figure 6 is the product $vev^2 e$. We are using juxtaposition here instead of $*$ to denote the product of cubes, since there is no other product under discussion on the set of cubes.

FIGURE 6. The product vev^2e

It is easy to check that the product $X \times X \to X$ induces the product previously defined on $F = \pi_1(X)$, if one is careful with basepoints. Namely, since $(v, v) \mapsto v^2$, we use the edge e to change basepoints so that there is an induced map on $\pi_1(X) = \pi_1(X, v)$.

Since X is a topological semigroup, we get a ring structure on $H_*(X) = H_*(F)$, which we are now in a position to calculate.

4. Calculation of the homology ring

Let $C = C(X)$ be the cellular chain complex of X with respect to its cubical structure. Then C is a differential graded ring (without identity). If we forget the differential, C is the (graded) ring of noncommuting polynomials without constant term, in two variables v, e with v in degree 0 and e in degree 1. The differential is determined by the formulas

$$\partial e = v^2 - v , \qquad \partial(xy) = \partial x \cdot y + (-1)^{\deg x} x \cdot \partial y .$$

Note, for the sake of intuition, that C is the universal example of a differential graded ring having an element v of degree 0 such that v^2 is homologous to v. This should be compared to the description of F in Section 1.2. We will use C to compute $H_*(F)$ as a ring. What makes the computation feasible is that, although C is a free abelian group of infinite rank in every dimension, it is "small" in the sense of having only two generators as a ring.

Let z be the commutator $[v, e] := ve - ev$. It is a 1-cycle, whose homology class $[z] \in H_1(C) = H_1(F)$ we denote by ζ. Let $\varepsilon = [v] \in H_0(C) = H_0(F) = \mathbb{Z}$. We will give two versions of the homology calculation. The first arises naturally from the method of proof, but the second exhibits the decomposition of the homology with respect to right and left multiplication by ε.

THEOREM 4.1.

(1) $H_*(F)$ is generated as a graded ring by ε and ζ, which satisfy the relations

$$\varepsilon^2 = \varepsilon$$
$$\varepsilon\zeta = \zeta - \zeta\varepsilon .$$

$H_n(F)$ is free abelian of rank 2 for all $n \geq 1$, with basis $\zeta^n, \zeta^n\varepsilon$.

(2) Let $\alpha = \zeta\varepsilon$ and $\beta = \varepsilon\zeta$. Then $H_*(F)$ is generated as a ring by ε, α, and β, which satisfy the relations

$$\varepsilon^2 = \varepsilon$$
$$\varepsilon\alpha = \beta\varepsilon = 0$$
$$\alpha\varepsilon = \alpha , \quad \varepsilon\beta = \beta .$$

We have $\alpha^2 = \beta^2 = 0$, and the alternating products $\alpha\beta\alpha\cdots$ and $\beta\alpha\beta\cdots$ form a basis for $H_*(F)$ in positive dimensions.

Proof. It will be convenient to adjoin an identity to C, getting a differential graded ring $R := \mathbb{Z} \oplus C$ with identity. It is the ring of all noncommuting polynomials in v, e or, equivalently, the tensor algebra over $\mathbb{Z}$ of the free $\mathbb{Z}$-module generated by v and e. We will compute $H_*(R)$, which is simply $H_*(F)$ with an extra summand $\mathbb{Z}$ in dimension 0.

Let $w = v^2 - v$ and let $S < R$ be the subring (with identity) generated by e, w, and z. I claim that S is the ring of noncommuting polynomials in the variables e, w, z and that $R = S \oplus Sv$. In other words, the set $\mathcal{M}$ consisting of the monomials in e, w, z and their products with v forms a basis for R as a $\mathbb{Z}$-module. We will show that $\mathcal{M}$ spans R; a counting argument that can be found in [20, p. 498] then shows that $\mathcal{M}$ is a basis. It suffices to show that any monomial in v, e can be rewritten as a linear combination of the monomials in $\mathcal{M}$. We may assume that v occurs in the given monomial and that the first occurence of v is not at the end. It is therefore followed either by another v or by e. If it is followed by v, then we replace the resulting v^2 by $w + v$; otherwise we replace ve by $ev + z$. We now have two terms, and we repeat the process with each of those. Continuing in this way, we arrive after finitely steps at a linear combination of elements of $\mathcal{M}$. (Termination of the process follows from the fact that we are moving v's to the right and, when we rewrite v^2, reducing the total number of v's.)

We now know that S is the tensor algebra $\mathbb{Z} \oplus D \oplus (D \otimes D) \oplus \cdots$, where D is the $\mathbb{Z}$-module generated by e, w, z. Each summand is a chain subcomplex, whose homology can be computed by the Künneth formula. By inspection, $H_*(D) = \mathbb{Z}$, concentrated in dimension 1 and generated by ζ; so we conclude that $H_n(S) = \mathbb{Z}$ for every $n \geq 0$, generated by ζ^n. The complementary summand Sv is isomorphic to S as a chain complex, so it contributes a second $\mathbb{Z}$, generated by $\zeta^n \varepsilon$, for each n. To complete the proof of (1), we need to check the relations. The relation $\varepsilon^2 = \varepsilon$ is immediate from $\partial e = v^2 - v$. For the second, one checks that $\partial(e^2)$ is the commutator $[v^2 - v, e] = [v^2, e] - z = vz + zv - z$, whence $\varepsilon \zeta + \zeta \varepsilon - \zeta = 0$.

Turning now to (2), the stated relations are immediate from the relations in (1). For example, we get $\varepsilon \zeta \varepsilon = 0$, which says precisely that $\varepsilon \alpha = \beta \varepsilon = 0$, by multiplying the relation $\varepsilon \zeta = \zeta - \zeta \varepsilon$ by ε on the right. The equation $\alpha^2 = 0$ now follows by computing $\alpha \varepsilon \alpha$ in two different ways, and similarly for $\beta^2 = 0$. Since $\zeta = \alpha + \beta$, it is clear that $H_*(R)$ is generated by $\varepsilon, \alpha, \beta$ and hence that the alternating products of α and β generate $H_n(R)$ additively for $n > 0$. Since we already know that $H_n(R) \cong \mathbb{Z} \oplus \mathbb{Z}$, we conclude that the alternating products form a basis. (Alternatively, note that $\zeta^n = (\alpha + \beta)^n$ is the sum of the two alternating products of length n, while $\zeta^n \varepsilon$ is the one that ends in α.) $\square$

Remark 4.2. In both (1) and (2), it is easy to show that the stated relations form a system of defining relations.

Remark 4.3. The subring $\varepsilon H_*(F) \varepsilon < H_*(F)$, on which ε is a 2-sided identity, is easily checked to be the polynomial ring $\mathbb{Z}[t]$ generated by $t := \beta \alpha \in H_2(F)$. I claim that it can be identified with $H_*(F')$, where F' is the commutator subgroup of F. To deduce this from what we have already done, recall first that F' is the subgroup of F consisting of dyadic PL-homeomorphisms with support in the interior of $[0, 1]$ (see Section 1.5). From this one sees that F' is closed under the product μ defined in Section 2.1, so $H_*(F')$ is a ring. Moreover, it is not hard to see that the canonical generator of $H_0(F')$ is a 2-sided identity, so the map $H_*(F') \to H_*(F)$ induced by the inclusion $F' \hookrightarrow F$ is a ring homomorphism with image contained

in $\mathbb{Z}[t]$. The next observation is that F behaves homologically as though it were $F' \times F_{ab}$; more precisely, there is a homomorphism $F \to F' \times F_{ab}$ that induces an isomorphism in homology (see [20, Lemma 4.1]). One can deduce that $H_*(F')$ maps injectively onto a direct summand of $H_*(F)$, and a counting argument [based on $H_*(F) \cong H_*(F') \otimes H_*(F_{ab})$] shows that the image is infinite cyclic in every even dimension. Since the image is a direct summand of $\mathbb{Z}[t]$, it must equal $\mathbb{Z}[t]$.

5. Calculation of the cup product

By standard arguments, as in the homology theory of topological groups, the diagonal map $F \to F \times F$ (or $X \to X \times X$) induces a ring homomorphism

$$\Delta \colon H_*(F) \to H_*(F) \otimes H_*(F) \;,$$

whose dual is the cup product in cohomology. Thus the cup product will be known if we can compute Δ on a set of ring generators of $H_*(F)$, such as $\varepsilon, \alpha, \beta$. By general principles, we know that $\Delta(\varepsilon) = \varepsilon \otimes \varepsilon$ and that α and β are primitive, i.e., $\Delta(\alpha) = \alpha \otimes \varepsilon + \varepsilon \otimes \alpha$ and similarly for β. One can deduce, after some calculations, that the cohomology ring is given by formula (0.1) in the introduction. The calculations are simpler if we make use of the homology equivalence $F \to F' \times F_{ab}$ mentioned in Remark 4.3 above, so we will phrase the statement and proof in those terms.

Recall that the divided polynomial ring $\Gamma(u)$ on one generator u is the subring of the polynomial ring $\mathbb{Q}[u]$ generated additively by the elements $u^{(i)} := u^i/i!$ $(i \geq 0)$. Note that

$$(5.1) \qquad\qquad u^{(i)} u^{(j)} = \binom{i+j}{i} u^{(i+j)} \;,$$

so the $\mathbb{Z}$-span of the $u^{(i)}$ is indeed a ring.

THEOREM 5.1.
  (1) $H^*(F') \cong \Gamma(u)$, where $\deg u = 2$.
  (2) There are ring isomorphisms

$$H^*(F) \cong H^*(F') \otimes H^*(F_{ab}) \cong \Gamma(u) \otimes \bigwedge(a, b) \;,$$

  where $\deg a = \deg b = 1$ and $\deg u = 2$.

PROOF. Recall that $H_*(F') = \mathbb{Z}[t]$. By general principles (or a calculation of $\Delta(t) = \Delta(\beta)\Delta(\alpha)$), the element $t \in H_2(F')$ is primitive, so

$$\Delta(t^n) = (t \otimes 1 + 1 \otimes t)^n = \sum_{i+j=n} \binom{i+j}{i} t^i \otimes t^j \;.$$

Dualizing, we obtain a basis $(u^{(n)})_{n \geq 0}$ for $H^*(F')$, with multiplication law as in (5.1). This proves (1), and (2) follows at once. $\square$

References

[1]   J. M. Belk, *Thompson's group F*, Ph.D. Thesis, Cornell University, 2004, `http://www.math.cornell.edu/~belk/Thesis.pdf`.
[2]   J. M. Belk and K. S. Brown, *Forest diagrams for elements of Thompson's group F*, Internat. J. Algebra Comput. (2004), to appear.
[3]   K. S. Brown, *Finiteness properties of groups*, J. Pure Appl. Algebra **44** (1987), no. 1-3, 45–75. MR 885095

[4]   ———, *The geometry of finitely presented infinite simple groups*, Algorithms and classification in combinatorial group theory (G. Baumslag, ed.), Mathematical Sciences Research Institute Publications, vol. 23, Springer-Verlag, New York, 1992. Papers from the Workshop on Algorithms, Word Problems and Classification in Combinatorial Group Theory held in Berkeley, California, January 1989, pp. 121–136. MR 1230631

[5]   ———, *Cohomology of groups*, Graduate Texts in Mathematics, vol. 87, Springer-Verlag, New York, 1994. Corrected reprint of the 1982 original. MR 1324339

[6]   K. S. Brown and R. Geoghegan, FP_∞ *groups and HNN extensions*, Bull. Amer. Math. Soc. (N.S.) **9** (1983), no. 2, 227–229. MR 707963

[7]   ———, *An infinite-dimensional torsion-free* FP_∞ *group*, Invent. Math. **77** (1984), no. 2, 367–381. MR 752825

[8]   ———, *Cohomology with free coefficients of the fundamental group of a graph of groups*, Comment. Math. Helv. **60** (1985), no. 1, 31–45. MR 787660

[9]   J. W. Cannon, W. J. Floyd, and W. R. Parry, *Introductory notes on Richard Thompson's groups*, Enseign. Math. (2) **42** (1996), no. 3-4, 215–256. MR 1426438

[10]   D. M. Clark, *Varieties with isomorphic free algebras*, Colloq. Math. **20** (1969), 181–187. MR 244128

[11]   J. Dydak, *A simple proof that pointed FANR-spaces are regular fundamental retracts of ANR's*, Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys. **25** (1977), no. 1, 55–62 (English, with Russian summary). MR 442918

[12]   ———, *1-movable continua need not be pointed 1-movable*, Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys. **25** (1977), no. 6, 559–562 (English, with Russian summary). MR 474226

[13]   P. Freyd and A. Heller, *Splitting homotopy idempotents. II*, J. Pure Appl. Algebra **89** (1993), no. 1-2, 93–106. MR 1239554

[14]   G. Higman, *Finitely presented infinite simple groups*, Department of Pure Mathematics, Department of Mathematics, I.A.S. Australian National University, Canberra, 1974. MR 376874

[15]   B. Jónsson and A. Tarski, *On two properties of free algebras*, Math. Scand. **9** (1961), 95–101. MR 126399

[16]   T. Napier and M. Ramachandran, *Thompson's group F is not Kähler*, Contemporary Mathematics, this volume.

[17]   D. Quillen, *Higher algebraic K-theory. I*, Algebraic K-theory. I: Higher K-theories, Springer-Verlag, Berlin, 1973. Edited by H. Bass; Lecture Notes in Mathematics, Vol. 341, pp. 85–147. MR 338129

[18]   ———, *Homotopy properties of the poset of nontrivial p-subgroups of a group*, Adv. in Math. **28** (1978), no. 2, 101–128. MR 493916

[19]   D. M. Smirnov, *Cantor algebras with one generator. I*, Algebra i Logika **10** (1971), 61–75 (Russian). MR 296006

[20]   M. Stein, *Groups of piecewise linear homeomorphisms*, Trans. Amer. Math. Soc. **332** (1992), no. 2, 477–514. MR 1094555

[21]   S. Świerczkowski, *On isomorphic free algebras*, Fund. Math. **50** (1961/1962), 35–44. MR 138573

DEPARTMENT OF MATHEMATICS, CORNELL UNIVERSITY, ITHACA, NY 14853

E-mail address: `kbrown@math.cornell.edu`