

Thompson’s group F and uniformly finite homology

DANIEL STALEY

We use the uniformly finite homology developed by Block and Weinberger to study the geometry of the Cayley graph of Thompson’s group F . In particular, a certain class of subgraph of F is shown to be nonamenable (in the Følner sense). This shows that if the Cayley graph of F is amenable, these subsets, which include every finitely generated submonoid of the positive monoid of F , must necessarily have measure zero.

20F65, 05C25; 43A07

1 Introduction

In 1965, Richard Thompson introduced his group F , which is finitely presented, has exponential growth, and contains subgroups isomorphic to $F \times F$. Every proper quotient of F is a quotient of $\mathbb{Z} \times \mathbb{Z}$. The question as to whether F is amenable was first posed in 1979. F is, in a sense, “on the edge of amenability”, as it is not elementary amenable but does not contain a free subgroup on two generators (see Brin and Squier [4]). If F is not amenable, it provides a finitely-presented counterexample to the Von Neumann conjecture. Very few such examples are known (Ol’shanskii and Sapir provided the first in 2000 [8]).

In 1955, Følner provided a geometric criterion for the amenability of a group based on the existence of subsets of the Cayley graph that satisfy a “small boundary” condition [6]. This criterion holds for semigroups as well (one may find a proof in Namioka [7]), and allows one to extend the definition of “amenable” to graphs of bounded degree. In 1992, Block and Weinberger extended the definition to a broad class of metric spaces [3]. They defined the *uniformly finite homology groups* $H_n^{\text{uf}}(M)$ of a metric space M and proved that M is amenable if and only if $H_0^{\text{uf}}(M) \neq 0$. This paper seeks to apply the results of Block and Weinberger to subgraphs of the Cayley graph of Thompson’s group F .

The main result of this paper is the following:

Theorem 1.1 Let k, l be nonnegative integers, with $l > 0$. Let Γ_k^l be the subgraph of the Cayley graph of F consisting of vertices that can be expressed in the form

$$a_1 \dots a_m b_1 \dots b_n,$$

where $m \leq k$, $a_1, \dots, a_m \in \{x_0, x_1, x_2, \dots\}$, and $b_1, \dots, b_n \in \{x_0, \dots, x_l\}$. Let the edge set of Γ_k^l include all edges in the Cayley graph of F that connect such vertices. Then Γ_k^l is not amenable.

The case $k = 1, l = 1$ was proved by D Savchuk in [9].

A corollary of this theorem is that all finitely-generated submonoids of the positive monoid of F are not amenable. It follows that if F is amenable these sets have measure zero.

Section 2 provides a very brief overview of Thompson's group F , and defines amenability and Følner's criterion. Readers interested in a more in-depth introduction are referred to Belk [1] or Canon and Floyd [5]. Section 3 discusses the results of Block and Weinberger and defines Ponzi flows. In Section 4 we prove the main result.

The author would like to thank Steve Ferry and the referee for numerous helpful comments and suggestions.

2 Thompson's group F

Thompson's group F can be described as the group with the following infinite presentation:

$$\langle x_0, x_1, x_2, x_3, \dots \mid x_j x_i = x_{i+1} x_j \text{ for } i > j \rangle$$

From this presentation, we see that $x_{i+1} = x_0 x_i x_0^{-1}$ for $i \geq 1$. Thus this group is finitely generated by $\{x_0, x_1\}$. F is finitely presented as well (see [1] or [5] for a proof). However, it is still useful to consider the infinite generating set $\{x_0, x_1, x_2, \dots\}$. We have the following definition:

Definition 2.1 The *positive monoid* of F is the submonoid of F consisting of elements that can be expressed as words in $\{x_0, x_1, x_2, \dots\}$ without using inverses.

Any element of F can be expressed as an element of the positive monoid multiplied by the inverse of such an element. Elements of F have a normal form which is such a product. The reader is referred to [1] or [5] for proofs.

In [1], the group F is studied using *two-way forest diagrams*. We will make extensive use of these diagrams when studying the positive monoid in Section 4. We describe

the two-way forest diagrams of the positive monoid here, referring the reader to [1] for the proofs.

Definition 2.2 A *binary forest* is an infinite sequence of binary trees, such that all but finitely many of the trees are trivial (ie have a single node):



Figure 1

Definition 2.3 A *pointed forest* is a binary forest with a distinguished, or “pointed”, tree:




Figure 2

For the remainder of this paper, we will omit the ellipses and assume a forest diagram or pointed forest diagram has an infinite number of trivial trees continuing to the right.

Each element of the positive monoid of F can be identified with a pointed forest. The identity element is the pointed forest consisting only of trivial trees, with the the pointer on the leftmost tree.

Right multiplication by x_0 moves the pointer one tree to the right (Figure 3).



Figure 3: Multiplication by x_0

Right multiplication by x_1 adds a caret between the pointed tree and the tree immediately to its right, making a new tree whose left child is the pointed tree and whose right child is the tree to its right. This new tree becomes the pointed tree (Figure 4).




Figure 4: Multiplication by x_1

Since $x_i = x_0^{i-1}x_1x_0^{-(i-1)}$, we can see that right multiplication by x_i moves the pointer $i - 1$ trees to the right, adds a caret, and then moves the pointer $i - 1$ trees to the left again. This is equivalent to adding a caret between the trees $i - 1$ and i steps to the right of the pointed tree (Figure 5).



Figure 5: Multiplication by x_3

Multiplication of pointed forests corresponds to “putting one on top of the other”: If P and Q are pointed forests, then PQ is the forest obtained by using the trees of P as the nodes of Q with the pointed tree in P attaching to the leftmost node of Q (Figure 6).




Figure 6: Multiplying the two pointed forests on the left yields the pointed forest on the right

The pointer is then placed above whatever tree was pointed in Q .

Definition 2.4 A group G is called *amenable* if there exists a *right-invariant measure* on G —a function μ that assigns to each subset $A \subset G$ a value $0 \leq \mu(A) \leq 1$ such that:

- (1) $\mu(G) = 1$.
- (2) μ is finitely additive: If A and B are disjoint subsets of G , then $\mu(A) + \mu(B) = \mu(A \cup B)$.
- (3) μ is G -invariant: For any $g \in G$ and any $A \subset G$, $\mu(A) = \mu(Ag)$.

A useful result for determining amenability is *Følner's Criterion*, which uses the Cayley graph of G . Recall that the Cayley graph of G is the graph obtained by taking a generating set S and using G as the vertex set, connecting two vertices g and g' by an edge if $g' = gs$ for some $s \in S$.

Theorem 2.5 (Følner's Criterion) *A group G is amenable if and only if, for any $\epsilon > 0$, there exists a finite subset A of vertices in the Cayley graph of G such that*

$$\frac{\#\partial(A)}{\#A} < \epsilon,$$

where $\#A$ is the number of vertices in A , and $\#\partial(A)$ is the number of edges connecting vertices of A to vertices outside A .

Følner's criterion can be applied to any graph of finite valence. In particular, we say such a graph is amenable if Følner's criterion holds for that graph. This allows us to state the following proposition:

Proposition 2.6 *Let Γ be the subgraph of the Cayley graph of Thompson's group F (using the x_0, x_1 generating set) consisting of vertices in the positive monoid of F and all edges between such vertices. Then Γ is amenable if and only if F is amenable.*

For a proof see Savchuk [9]. The essential fact is that any finite subset of F can be translated into the positive monoid.

3 Uniformly finite homology

This section describes the uniformly finite homology of Block and Weinberger defined in [3]. We will always be considering a graph Γ of bounded degree, though many of their results apply to a much broader class of metric spaces.

Definition 3.1 Let Γ be a graph of bounded degree with vertex set V . A *uniformly finite 1-chain with integer coefficients* on Γ is a formal infinite sum $\sum a_{x,y}(x, y)$, where the (x, y) are ordered pairs of vertices of Γ , $a_{x,y} \in \mathbb{Z}$, such that the following properties are satisfied:

- (1) There exists $K > 0$ such that $|a_{x,y}| < K$ for all vertices x and y .
- (2) There exists $R > 0$ such that $a_{x,y} = 0$ whenever $d(x, y) > R$.

Notice that condition (2) guarantees that for any fixed $x \in V$, the set of pairs (x, y) such that $a_{x,y} \neq 0$ is finite. This allows us to make the following definition:

Definition 3.2 A uniformly finite 1–chain is a *Ponzi scheme* if, for all $x \in \Gamma$, we have $\sum_{v \in \Gamma} a_{v,x} - \sum_{v \in \Gamma} a_{x,v} > 0$.

We now state the main result of [3] that we will use in this paper:

Theorem 3.3 *Let Γ be a graph of bounded degree. A Ponzi scheme exists on Γ if and only if Γ is not amenable.*

We will use a rephrased version of this theorem for the case of our graphs:

Definition 3.4 Let Γ be a graph of bounded degree with vertex set V . A *Ponzi flow* on Γ will mean a function $\Phi: V \times V \rightarrow \mathbb{Z}$ with the following properties:

- (i) $\Phi(a, b) = 0$ if there is no edge from a to b in Γ ,
- (ii) $\Phi(a, b) = -\Phi(b, a)$ for all $a, b \in V$.
- (iii) The function Φ is bounded.
- (iv) For each $a \in V$, $\sum_{b \in V} \Phi(b, a) > 0$.

Note that the sum in condition (iv) is guaranteed to be finite by condition (i). A Ponzi flow is almost exactly a Ponzi scheme in different language, with the exception that all “pairs” must be of distance 1. However, this difference is unimportant:

Proposition 3.5 *Let Γ be a graph of bounded degree. There exists a Ponzi scheme on Γ if and only if there exists a Ponzi flow on Γ .*

Proof The “if” direction is trivial: Given a Ponzi flow, we simply define our formal sum to be $\sum \Phi(x, y)(x, y)$. This will be a uniformly finite 1–chain with integer coefficients, as condition (1) is implied by (iii), and condition (2) is implied by (i). This 1–chain will be a Ponzi scheme by conditions (ii) and (iv).

To see the “only if” direction, one can take the coefficients $a_{x,y}$ with $d(x, y) > 1$ and “reroute the flow” by replacing them with a sequence of distance-one steps along a shortest path from x to y . Since the graph has bounded degree and the coefficients $a_{x,y}$ are bounded, the result is still a Ponzi scheme, with $a_{x,y}$ nonzero only when $d(x, y) = 1$. This is easily converted into a Ponzi flow by setting $\Phi(x, y) = a_{x,y} - a_{y,x}$.

This completes the proof. \square

A quantitative treatment of Proposition 3.5 and of Ponzi flows can be found in Benjamini, Lyons, Perez and Schramm [2].

If a Ponzi flow exists on a Cayley graph, there can be no right-invariant measure on the group since the group cannot be amenable. Indeed, the following proposition follows from the results of Block and Weinberger:

Proposition 3.6 *If a group Γ is amenable but its Cayley graph contains a nonamenable subgraph P , then for any right-invariant measure μ on Γ , $\mu(P) = 0$.*

4 Large nonamenable subgraphs of F

In this section we will prove Theorem 1.1.

We begin by characterizing the two-way forest diagrams of Γ_k^l . Given any binary tree T on n nodes, we define $s(T)$ to be the forest obtained by removing all the carets along the left edge of T (Figure 7).




Figure 7: Applying s to a tree removes the left carets as shown

We extend the definition of s to apply to forests, as well as single trees, by applying s separately to each nontrivial tree in the forest. We will define the *complexity* of a tree or forest to be the minimum number of applications of s required to turn it into a forest of only trivial trees.

Note that applying s to a tree T gives a forest whose rightmost tree is the right child of T , and the remainder of the forest is s applied to the left child of T . This gives the following:

Proposition 4.1 *The complexity of a tree is the maximum of the complexity of its left child and one more than the complexity of its right child.*

We record here two basic properties of complexity and the function s :

Proposition 4.2 *Let T be a tree on n nodes, and let R be an n -tree forest. Denote by RT the tree obtained by attaching the roots of R to the nodes of T . If T has complexity j , then $s^j(RT)$ consists only of carets in R , ie every caret from T is removed by s^j . (Figure 8.)*

Proof This is easy to see, as we can determine whether a caret is removed by s^j by examining its relationship with those above it. Namely, when we examine the unique path from a caret to the root of the tree, it consists of moves to the right and moves




Figure 8: A 4-node tree T of complexity 2(top left) and a 4-tree forest R (bottom left) multiply to give RT (top right), and each caret of T is removed in $s^2(RT)$ (bottom right)

to the left. An application of s removes all carets whose path consists only of moves to the right. Further, any caret’s path to the root hits the left edge at some point, and consists only of moves to the right afterwards. After s is applied, the path is truncated, starting from the move that reaches the left edge (which is a move to the left). Thus each new path from a remaining caret to the root of its new tree is left with one less move to the left after applying s . So s^j removes all carets whose paths contain $j - 1$ or fewer moves to the left. Since this property is unchanged in the carets of T whether or not it sits on R , the effect of s^j is the same on carets of T , ie it removes them all. □

Proposition 4.3 *A pointed forest diagram consisting of a single nontrivial tree T of complexity j in the leftmost position, with the pointer on that tree, can be expressed as word in x_1, \dots, x_j .*

Proof We will proceed by induction on the number of carets in the tree T . Suppose the statement is true for all trees with n or fewer carets, and let T be a tree with $n + 1$ carets and complexity j . Then the left child of T has no more than n carets and complexity no more than j by Proposition 4.1. Thus by the inductive hypothesis the left child can be constructed as a word w in x_1, \dots, x_j . The right child of T has no more than n carets and complexity no more than $j - 1$ by Proposition 4.1, thus can be constructed as a word v in x_1, \dots, x_{j-1} . We can construct the desired pointed forest as $wx_0vx_0^{-1}x_1$, since this will construct the left child, move the pointer to the right, construct the right child, move the pointer back to the left child, and finally construct the top caret. However, since $x_i = x_0x_{i-1}x_0^{-1}$, by inserting $x_0^{-1}x_0$ between each letter of v we can rewrite $x_0vx_0^{-1}$ as a word in x_2, \dots, x_j . Thus, the word $wx_0vx_0^{-1}x_1$ can be rewritten as a word in x_1, \dots, x_j , and the proposition is proved. □

For a positive integer j , we define a function ϕ_j from pointed forests to forests in the following way: Apply s^j to the pointed tree and every tree to its left. For each positive

integer $q < j$, apply s^{j-q} to the tree that is q trees to the right of the pointed tree. That is, apply s^{j-1} to the tree to the immediate right of the pointed tree, s^{j-2} to next tree to the right, etc.

For the proof of the main theorem we will use the following lemma. Recall that Γ_k^l is the subgraph of the Cayley graph of F consisting of vertices that can be expressed in the form $a_1 \dots a_m b_1 \dots b_n$, with $m \leq k$, $a_1, \dots, a_m \in \{x_0, x_1, x_2, \dots\}$, and $b_1, \dots, b_n \in \{x_0, \dots, x_l\}$.

Lemma 4.4 *A pointed forest P lies in Γ_k^l if and only if $\phi_l(P)$ has k or fewer carets.*

Proof Let $P \in \Gamma_k^l$. First suppose that $k = 0$. In this case P can be expressed as a word v in $\{x_0, \dots, x_l\}$, and the proposition says it is annihilated by ϕ_l , ie $\phi_l(P)$ consists only of trivial trees. We proceed by induction on the length of v . If $\phi_l(v)$ consists only of trivial trees, then so does $\phi_l(vx_0)$, since $\phi_l(vx_0)$ is a subforest of $\phi_l(v)$ (each tree has s applied to it either the same number of times or one more time, since the pointer has simply moved one tree to the right).

For $0 < i \leq l$, multiplying by x_i adds a caret to the right of the tree $i - 1$ trees from the pointer, combining it with the next tree to make a new tree. Since the left and right children of this new tree were $i - 1$ and i trees to the right of the caret, respectively, by induction their respective complexities are no more than $(l - (i - 1))$ and $(l - i)$. Thus by Proposition 4.1 the new tree has complexity no more than $(l - i + 1)$. Since this new tree is $i - 1$ trees to the right of the pointer, it is still annihilated by ϕ_l . The trees to the left of the new caret are unchanged, and the trees to the right of the caret have each been brought 1 tree closer to the pointer since two of the intervening trees have been merged. Thus ϕ_l applies s an additional time to each of these trees. This means that they will still be annihilated by ϕ_l , and so the new pointed forest is still turned into a trivial forest by ϕ_l .

The above argument shows that $\phi_l(v)$ is trivial if $v \in \Gamma_0^l$. Now let $P = wv$, where $w = a_1 \dots a_m$ with $a_i \in \{x_0, x_1, \dots\}$ and $m \leq k$ as in Theorem 1.1. Then wv attaches the trees of w to the nodes of v . Thus all the carets added in each tree of v are still removed by ϕ_l by Proposition 4.2, since s is applied the same number of times to each tree. Thus $\phi_l(wv)$ has at most the same number of carets as w , ie k or fewer. This proves the "only if" direction of the Lemma.

To prove the reverse direction, suppose that P is a pointed forest such that $\phi_l(P)$ has k or fewer carets. We can then create w as above to put these carets in place without moving the pointer (the generator x_i creates a caret on the i^{th} tree without moving the pointer).

Consider the element $w^{-1}P$. This is the pointed forest obtained by taking the trees in P that remain after applying ϕ_l , and replacing them with trivial trees (Figure 9).




Figure 9: For $l = 2$, if P is the forest in the top left then w is $\phi_2(P)$ with the pointer on the first tree (bottom left), and $w^{-1}P$ is shown on the right.

The resulting pointed forest is then annihilated by ϕ_l by Proposition 4.2, and so each tree under or to the left of the pointer has complexity at most l . Thus, we may construct these trees as words in x_1, \dots, x_l using Proposition 4.3 and inserting x_0 between each word. This will result in building the first tree, moving the pointer to the right, building the next tree, etc. Further, the tree that is j trees to the right of the pointer has complexity at most $l - j$, and so Proposition 4.3 says we can construct it as $x_0^j u x_0^{-j}$, where u is a word in x_1, \dots, x_{l-j} . As above, we then insert $x_0^{-j} x_0^j$ between each letter of u , which allows us to rewrite it as a word in x_{j+1}, \dots, x_l . Repeating this for each j and appending these words in increasing order constructs all trees to the right of the pointer. This completes the construction of $w^{-1}P$ as a word in x_0, \dots, x_l , which we will call v . Thus, $P = w w^{-1}P = wv$, and the proof is complete. \square

We are now ready to prove the main theorem, which will occupy the remainder of this section.

Proof of Theorem 1.1 Let $P \in \Gamma_k^l$. Note that applying ϕ_l to P affects at most l trees under or to the right of the pointer. Thus, by Lemma 4.4 there are at most $k + l$ nontrivial trees under or to the right of the pointer in P , otherwise, $\phi_l(P)$ would have more than k nontrivial trees and thus certainly have more than k carets.

For any $P \in \Gamma_k^l$, let $c(P)$ be the number of nontrivial trees T to the left of the pointed tree such that if the pointer is moved to T the resulting pointed forest remains in Γ_k^l . By the preceding paragraph $c(P)$ is never more than $k + l$. Now define $\Phi: \Gamma_k^l \times \Gamma_k^l \rightarrow \mathbb{Z}$ as follows: For each pointed forest P ,

- If $Px_0^{-1} \in \Gamma_k^l$, set $\Phi(P, Px_0^{-1}) = c(P)$ and $\Phi(Px_0^{-1}, P) = -c(P)$.
- If the pointed tree in P is nontrivial, and $Px_1^{-1} \in \Gamma_k^l$, set $\Phi(P, Px_1^{-1}) = 1$ and $\Phi(Px_1^{-1}, P) = -1$.

- For all other pairs (P, P') , set $\Phi(P, P') = 0$.

We claim that Φ is a Ponzi flow.

It is clear from the definition that Φ satisfies conditions (i) and (ii) in Definition 3.4, and since $c(P) \leq k + l$ for each P , Φ also satisfies condition (iii). It thus remains only to check condition (iv). So we shall consider a pointed forest $P \in \Gamma_k^l$.

For ease of notation we will use the convention $\Phi(P', P) = 0$ if either P' or P is not in Γ_k^l . Since each vertex has 4 neighbors in the Cayley graph, this lets us state

$$\sum_{P' \in \Gamma_k^l} \Phi(P', P) = \Phi(Px_0, P) + \Phi(Px_0^{-1}, P) + \Phi(Px_1, P) + \Phi(Px_1^{-1}, P).$$

We also have

$$\begin{aligned} \Phi(Px_1, P) + \Phi(Px_1^{-1}, P) &\geq 1 - 1 = 0, \\ \Phi(Px_0, P) + \Phi(Px_0^{-1}, P) &\geq c(Px_0) - c(P) \geq 0. \end{aligned}$$

The first inequality holds since Px_1 is always in Γ_k^l , and the second holds since any tree to the left of the pointed tree in P is also to the left of the pointed tree in Px_0 , so $c(P) \leq c(Px_0)$. Further, since the only tree counted in $c(Px_0)$ but not in $c(P)$ is the pointed tree in P , we have that $\Phi(Px_0, P) + \Phi(Px_0^{-1}, P) > 0$ exactly when the pointed tree in P is nontrivial. Thus, if the pointed tree is nontrivial, condition (iv) is satisfied since

$$\sum_{P' \in \Gamma_k^l} \Phi(P', P) \geq \Phi(Px_0, P) + \Phi(Px_0^{-1}, P) > 0.$$

If the pointed tree in P is trivial, Px_1^{-1} is not in Γ_k^l , so $\Phi(Px_1, P) = 1$ and $\Phi(Px_1^{-1}, P) = 0$, and thus

$$\sum_{P' \in \Gamma_k^l} \Phi(P', P) = \Phi(Px_0, P) + \Phi(Px_0^{-1}, P) + \Phi(Px_1, P) \geq 1$$

Thus condition (iv) holds for all pointed forests where the pointed tree is trivial as well, and so Φ is a Ponzi flow. □

We close with some immediate corollaries:

Corollary 4.5 *If F is amenable, then for any right-invariant measure μ , $\mu(\Gamma_k^l) = 0$.*

Proof By Theorem 1.1, Γ_k^l has a Ponzi flow, and thus by Proposition 3.6 it always has measure zero. □

Corollary 4.6 *If F is amenable, then for any right-invariant measure μ , and any finitely generated submonoid M of the positive monoid, $\mu(M) = 0$.*

Proof Letting p_1, \dots, p_n be generators of M , express each as a word in the generating set $\{x_0, x_1, x_2, \dots\}$. Let L be the maximum index of the x_i used to express the p_j ; then M is a subset of the monoid generated by x_0, x_1, \dots, x_L . But this monoid is exactly Γ_0^L , which by the previous corollary has measure zero. Thus, $\mu(M) = 0$. \square

References

- [1] **J Belk**, *Thompson's Group F* , PhD thesis, Cornell University (2004)
- [2] **I Benjamini, R Lyons, Y Peres, O Schramm**, *Group-invariant percolation on graphs*, *Geom. Funct. Anal.* 9 (1999) 29–66 MR1675890
- [3] **J Block, S Weinberger**, *Aperiodic tilings, positive scalar curvature and amenability of spaces*, *J. Amer. Math. Soc.* 5 (1992) 907–918 MR1145337
- [4] **M G Brin, C C Squier**, *Groups of piecewise linear homeomorphisms of the real line*, *Invent. Math.* 79 (1985) 485–498 MR782231
- [5] **J W Cannon, W J Floyd, W R Parry**, *Introductory notes on Richard Thompson's groups*, *Enseign. Math.* (2) 42 (1996) 215–256 MR1426438
- [6] **E Følner**, *On groups with full Banach mean value*, *Math. Scand.* 3 (1955) 243–254 MR0079220
- [7] **I Namioka**, *Følner's conditions for amenable semi-groups*, *Math. Scand.* 15 (1964) 18–28 MR0180832
- [8] **A Y Ol'shanskii, M V Sapir**, *Non-amenable finitely presented torsion-by-cyclic groups*, *Publ. Math. Inst. Hautes Études Sci.* (2002) 43–169 (2003) MR1985031
- [9] **D Savchuk**, *Some graphs related to Thompson's group F* arXiv:0707.4316

*Department of Mathematics, Rutgers University
110 Frelinghuysen Road, Piscataway NJ 08854-8019, USA*

staley@math.rutgers.edu

math.rutgers.edu/~staley

Received: 3 July 2008 Revised: 1 September 2009