

Bowdoin College

Bowdoin Digital Commons

Mathematics Faculty Publications

Faculty Scholarship and Creative Work

12-1-2003

Thompson's group F is not almost convex

Sean Cleary

City College of New York

Jennifer Taback

State University of New York Albany

Follow this and additional works at: <https://digitalcommons.bowdoin.edu/mathematics-faculty-publications>

Recommended Citation

Cleary, Sean and Taback, Jennifer, "Thompson's group F is not almost convex" (2003). *Mathematics Faculty Publications*. 63.

<https://digitalcommons.bowdoin.edu/mathematics-faculty-publications/63>

This Article is brought to you for free and open access by the Faculty Scholarship and Creative Work at Bowdoin Digital Commons. It has been accepted for inclusion in Mathematics Faculty Publications by an authorized administrator of Bowdoin Digital Commons. For more information, please contact mdoyle@bowdoin.edu, a.sauer@bowdoin.edu.



Thompson's group F is not almost convex

Sean Cleary^{a,1,*} and Jennifer Taback^{b,2}

^a *Department of Mathematics, City College of New York, City University of New York, New York, NY 10031, USA*

^b *Department of Mathematics and Statistics, University at Albany, Albany, NY 12222, USA*

Received 1 September 2002

Communicated by Efim Zelmanov

Abstract

We show that Thompson's group F does not satisfy Cannon's almost convexity condition $AC(n)$ for any positive integer n with respect to the standard generating set with two elements. To accomplish this, we construct a family of pairs of elements at distance n from the identity and distance 2 from each other, which are not connected by a path lying inside the n -ball of length less than k for increasingly large k . Our techniques rely upon Fordham's method for calculating the length of a word in F and upon an analysis of the generators' geometric actions on the tree pair diagrams representing elements of F .

© 2003 Elsevier Inc. All rights reserved.

Keywords: Thompson's group F ; Almost convexity; Tree pair diagrams

1. Introduction

Cannon [7] introduced the notion of almost convexity for a group G with respect to a finite generating set X . This finite generating set X determines a word metric d_X for G and its Cayley graph. G is *almost convex* (k) or $AC(k)$ with respect to X if there is a number $N(k)$ so that for all positive integers n , given two elements y and z in the ball $B(n)$ of radius n with $d_X(y, z) = k$, there is a path γ from y to z of length at most $N(k)$ which lies entirely in $B(n)$. Cannon showed that if a group G is $AC(2)$ with respect to a finite generating set then G is $AC(k)$ for $k \geq 2$ and thus a group satisfying $AC(2)$ is called

* Corresponding author.

E-mail addresses: cleary@sci.ccny.cuny.edu (S. Cleary), jtaback@math.albany.edu (J. Taback).

¹ The author acknowledges support from PSC-CUNY grant #63438-0032.

² The author thanks the University of Utah for their hospitality during the writing of this paper.

almost convex with respect to that generating set. If a group is almost convex with respect to any finite generating set, we say the group is *almost convex*. Almost convexity allows algorithmic construction of $B(n+1)$ from $B(n)$ by making it sufficient to consider only a finite set of possible ways that an element in $B(n+1)$ can be obtained from different elements of $B(n)$.

A number of families of groups have been shown to be almost convex. Cannon [7] showed that hyperbolic groups are almost convex and that amalgamated products of almost convex groups are almost convex. Stein and Shapiro [13] showed that fundamental groups of closed three manifolds whose geometry is not modeled on *Sol* are almost convex. Other families of groups have been shown not to be almost convex. Cannon, Floyd, Grayson and Thurston [5] showed that fundamental groups of manifolds with *Sol* geometry are not almost convex, and Miller and Shapiro [12] showed that the solvable Baumslag–Solitar groups $BS(1, n)$ are not almost convex. Unfortunately, the property of almost convexity can depend upon generating set. Thiel [14] showed that generalized Heisenberg groups are not almost convex with respect to the generating sets in their standard presentations, but are almost convex with respect to some finite generating sets from alternate presentations.

Although Thompson's group F has been studied extensively in many branches of mathematics, the metric properties of F were poorly understood until recently. Burillo [4] and Burillo, Cleary and Stein [3] developed estimates for measuring distance in F , and Fordham [8] developed a remarkable method for computing distance in F .

We prove below that F does not satisfy Cannon's $AC(2)$ property in its standard finite generating set, and thus is not almost convex with respect to that generating set.

Thompson's group F has a number of different manifestations. Originally discovered by Thompson [15], in logic F is understood as the group of automorphisms of a free algebra. F also has connections with homotopy theory developed by Freyd and Heller [9,10], groups of homeomorphisms of the interval studied by Brin and Squier [1] and Brown and Geoghegan [2] and diagram groups defined by Guba and Sapir [11]. Cannon, Floyd and Parry [6] give an introduction to and summarize many of the remarkable properties of F .

Thompson's group F has the infinite presentation \mathcal{P} given by

$$\mathcal{P} = \langle x_k, k \geq 0 \mid x_i^{-1} x_j x_i = x_{j+1} \text{ if } i < j \rangle.$$

We can see that the lower index generators conjugate the higher-index generators by incrementing their indices. Since x_0 conjugates x_1 to x_2 and successively to all higher index generators, it is clear that F is finitely generated. In fact, all of the infinitely many relators in \mathcal{P} are consequences of a basic set of two relators. Thus, there is the following standard finite presentation \mathcal{F} for F :

$$\mathcal{F} = \langle x_0, x_1 \mid x_0 x_1^{-1}, x_0^{-1} x_1 x_0, x_0 x_1^{-1}, x_0^{-2} x_1 x_0^2 \rangle.$$

We prove the following theorem:

Theorem 1.1. *Thompson's group F does not satisfy Cannon's almost convexity condition $AC(2)$ with respect to the generators in the standard finite presentation \mathcal{F} for F .*

We immediately obtain the corollary:

Corollary 1.2. *Thompson’s group F does not satisfy Cannon’s almost convexity condition $AC(n)$ for any positive integer $n > 2$ with respect to the generators in the standard finite presentation for F .*

In all of the following, we will consider the convexity properties of F only with respect to the standard generating set of two generators x_0 and x_1 .

2. Background on

Analytically, we define F as the group of orientation-preserving piecewise-linear homeomorphisms from $[0, 1]$ to itself where each homeomorphism has only finitely many singularities of slope, all such singularities lie in the dyadic rationals $\mathbf{Z}[\frac{1}{2}]$, and, away from the singularities, the slopes are powers of 2.

Combinatorially, F has the infinite and finite presentations given above. There is a convenient set of normal forms for elements of F in the infinite presentation \mathcal{P} given by $x_{i_1}^{r_1} x_{i_2}^{r_2} \cdots x_{i_k}^{r_k} x_{j_l}^{s_l} \cdots x_{j_2}^{s_2} x_{j_1}^{s_1}$ with $r_i, s_i > 0$, $i_1 < i_2 < \cdots < i_k$ and $j_1 < j_2 < \cdots < j_l$. This normal form is unique if we further require that when both x_i and x_i^{-1} occur, so does x_{i+1} or x_{i+1}^{-1} , as discussed by Brown and Geoghegan [2]. In what follows, when we refer to a word in normal form, we always mean the unique normal form.

The geometric description of F is in terms of tree pair diagrams. A tree pair diagram is a pair of rooted binary trees with the same number of leaves, as described in [6]. We number the leaves of each tree from left to right, beginning with 0. We refer to an interior node together with the two downward-directed edges from the node as a *caret*. We define the right (respectively left) child of a caret C to be the caret C_R (respectively C_L) which is attached to the right (left) downward edge of caret C .

Each tree in a tree pair can be regarded as a set of instructions for successive subdivision of the unit interval: the root caret subdivides the interval in half, a right child of the root subdivides $[\frac{1}{2}, 1]$ in half, and so on. This gives a correspondence between elements of F in the geometric description and the analytic description as follows. Let (T^-, T_+) be a pair of trees each with n leaves. Each tree determines a subdivision of $[0, 1]$ into n subintervals. The tree pair (T^-, T_+) corresponds to the piecewise linear homomorphism which maps the subintervals of the T^- subdivision to the subintervals of the T_+ subdivision, in order. This equivalence and the group operation are described in [6]. We refer to T^- as the *negative tree* and T_+ as the *positive tree* of the pair (T^-, T_+) .

A tree pair diagram is *unreduced* if each of T^- and T_+ contain a caret with leaves numbered m and $m + 1$, and it is *reduced* otherwise. Note that there are many tree pair diagrams representing the same element of F but there is a unique reduced tree pair diagram for each element of F . When we write (T^-, T_+) to represent an element of F , we are assuming that the tree pair is reduced.

If $x = (T^-, T_+)$ is a reduced pair of trees representing x , the normal form for x can be constructed by the following process, described in [6]. Beginning with the tree pair

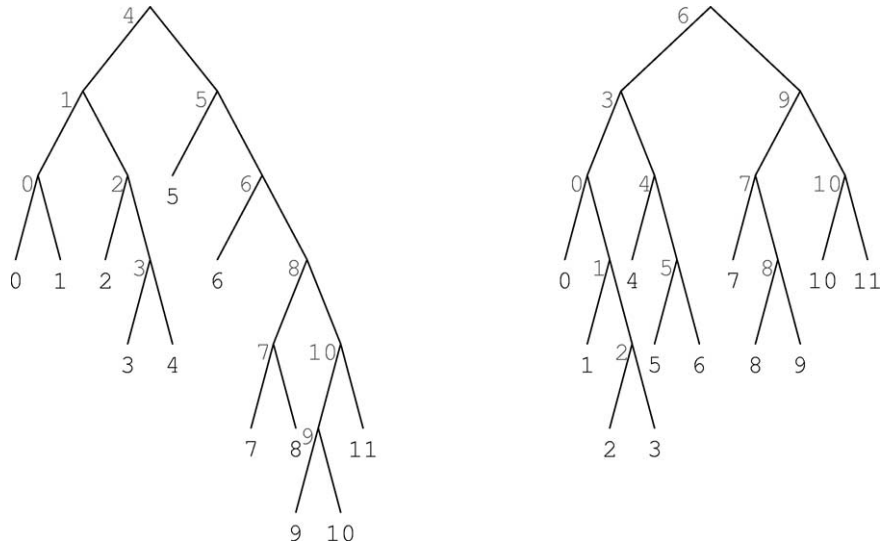


Fig. 1. Tree pair diagram for $x_0^2 x_1 x_2 x_4 x_5 x_7 x_8 x_9^{-1} x_7^{-1} x_3^{-1} x_2^{-1} x_0^{-2}$ with carets and leaves numbered.

(T^-, T_+) , we number the leaves of T^- and T_+ from left to right, beginning with 0. The exponent of the leaf labelled n , written $E(n)$, is defined as the length of the maximal path consisting entirely of left edges from n which does not reach the right side of the tree. Note that $E(n) = 0$ for a leaf labelled n which is a right child of a caret, as there is no path consisting entirely of left edges originating from n .

We compute $E(n)$ for all leaves in T^- , numbered 0 through m . The negative part of the normal form for x is then $x_m^{E(m)} x_{m-1}^{E(m-1)} \dots x_1^{E(1)} x_0^{E(0)}$. We compute the exponents for the leaves of the positive tree and thus obtain the positive part of the normal form as $x_0^{E(0)} x_1^{E(1)} \dots x_m^{E(m)}$. Many of the exponents may be 0, and after deleting these, we can index the remaining terms to correspond to the normal form given above, as detailed in [6].

In the tree pair diagram in Fig. 1, the exponent $E(0)$ of the leaf labelled 0 of T^- is 2 since there is a path of two left edges from leaf 0 which does not reach the right hand side of the tree. The third left edge emanating from leaf 0 touches the right-hand side of the tree and thus does not contribute to the exponent. The exponents of all the leaves of T^- are, in order, 2, 0, 1, 1, 0, 0, 0, 1, 0, 1, 0, 0, and the exponents of the leaves of T_+ are, in order, 2, 1, 1, 0, 1, 1, 0, 1, 1, 0, 0, 0. Using these exponents, and omitting any which are 0, we see that the tree pair diagram of Fig. 1 represents the word $x_0^2 x_1 x_2 x_4 x_5 x_7 x_8 x_9^{-1} x_7^{-1} x_3^{-1} x_2^{-1} x_0^{-2}$, in normal form.

If R is a caret on the right side of the tree with a single left leaf labeled k , then $E(k) = 0$ by definition. We use this fact to show that without loss of generality, T^- and T_+ may be assumed to have the same number of carets.

Suppose that T^- has k fewer carets than T_+ , and let the rightmost leaf of T^- be numbered m . Attach a single caret to leaf m in T^- , obtaining a new tree T' . It is easily computed that in T' , the final two exponents, $E(m)$ and $E(m + 1)$, are both 0. Thus the element of F represented by the tree pair (T', T_+) is identical to the element represented

by the tree pair (T^-, T_+) , and T' has one more caret than T^- . Repeat this process $k - 1$ additional times, with each repetition adding a caret to the rightmost leaf of the negative tree. This has no effect on the normal form of the resulting element, and increases the number of carets in the negative tree of the pair. Thus without loss of generality we may assume that T^- and T_+ have the same number of carets.

Similarly, given an element x in normal form with respect to the infinite generating set, it is possible to construct a tree pair diagram (T^-, T_+) so that each leaf has the correct exponent. In particular, the number of left edges of T^- emanating from the root caret is one more than the exponent of x_0^{-1} in the normal form and the number of left edges of T_+ emanating from the root caret is one more than the exponent of x_0 in the normal form for x .

The processes described above relate the normal form of words in F in the infinite presentation \mathcal{P} to the tree pair representation. For many questions involving the geometry of F , we must consider the length of words in F with respect to a metric arising from a finite generating set. Burillo [4] presented a way of estimating the word length $|x|_{\mathcal{F}}$ in the finite generating set \mathcal{F} from the normal form, which was refined by Burillo, Cleary, and Stein in [3].

Theorem 2.1 (Burillo [4, Proposition 2]; Burillo, Cleary, and Stein [3, Theorem 1]). *Let $w \in F$ have normal form $w = x_{i_1}^{r_1} \cdots x_{i_n}^{r_n} x_{j_m}^{s_m} \cdots x_{j_1}^{s_1}$, and let $D(w) = r_1 + r_2 + \cdots + r_n + s_1 + s_2 + \cdots + s_m + i_n + j_m$. Then*

$$\frac{D(w)}{3} \leq |w|_{\mathcal{F}} \leq 3D(w).$$

Burillo, Cleary, and Stein [3] also estimated of the length $|w|_{\mathcal{F}}$ of a word w given by a tree pair diagram in terms of the number of carets $N(w)$ in either tree.

2.1. Fordham's method of calculating word length

Fordham [8] presents a method of calculating the exact word length in F given a reduced pair of trees representing an element $x \in F$. We make some preliminary definitions before explaining Fordham's technique.

Let T be a finite rooted binary tree. The *left side* of T is the maximal path of left edges beginning at the root of T . Similarly, we have the *right side* of T . A caret in T is a *left caret* if its left edge is on the left side of the tree, a *right caret* if it is not the root and its right edge is on the right side of the tree, and an *interior caret* otherwise. The carets in T are numbered according to the infix ordering of nodes. We begin numbering with leaf 0 as the leftmost leaf and caret 0 the left caret whose left child is leaf 0. We number the left children of a caret before the caret itself, and number the right children after numbering the caret. The trees in Fig. 1 have their carets numbered according to this method.

Fordham classifies carets into seven disjoint types:

- (1) L_0 : The first caret on the left side of the tree, with caret number 0. Every tree has exactly one caret of type L_0 .

- (2) L_L : Any left caret other than the one numbered 0.
- (3) I_0 : An interior caret which has no right child.
- (4) I_R : An interior caret which has a right child.
- (5) R_I : Any right caret numbered k with the property that caret $k + 1$ is an interior caret.
- (6) R_{NI} : A right caret which is not an R_I but for which there is a higher numbered interior caret.
- (7) R_0 : A right caret with no higher-numbered interior carets.

The root caret is always considered to be a left caret of type L_L unless it has no left children, in which case it is the L_0 caret.

Working from caret 0 to caret 10, in infix order, in the tree T from Fig. 1, we see that the carets are of types

$$L_0, L_L, I_R, I_0, L_L, R_{NI}, R_I, I_0, R_I, I_0, \text{ and } R_0.$$

The carets in the tree T_+ of Fig. 1, in infix order, are of types

$$L_0, I_R, I_0, L_L, I_R, I_0, L_L, I_R, I_0, R_0, \text{ and } R_0.$$

The main result of Fordham [8] is that the word length $|x|_F$ of $x = (T, T_+)$ can be computed from knowing the caret types of the carets in the two trees, as long as they form a reduced pair, via the following process. We number the $k + 1$ carets according to the infix method described above, and for each i with $0 \leq i \leq k$ we form the pair of caret types consisting of the type of caret number i in T and the type of caret number i in T_+ . The single caret of type L_0 in T will be paired with the single caret of type L_0 in T_+ , and for that pairing we assign a weight of 0. For all other caret pairings, we assign weights according to the following symmetric table:

| | R_0 | R_{NI} | R_I | L_L | I_0 | I_R |
|----------|-------|----------|-------|-------|-------|-------|
| R_0 | 0 | 2 | 2 | 1 | 1 | 3 |
| R_{NI} | 2 | 2 | 2 | 1 | 1 | 3 |
| R_I | 2 | 2 | 2 | 1 | 3 | 3 |
| L_L | 1 | 1 | 1 | 2 | 2 | 2 |
| I_0 | 1 | 1 | 3 | 2 | 2 | 4 |
| I_R | 3 | 3 | 3 | 2 | 4 | 4 |

Fordham’s remarkable result is that the sum of these weights is exactly the length of the word in the word metric arising from the finite generating set.

Theorem 2.2 (Fordham [8, Theorem 2.5.1]). *Given a word $w \in F$ described by the reduced tree pair diagram (T, T_+) , the length $|w|_F$ of the word with respect to the generating set \mathcal{F} is the sum of the weights of the caret pairings in (T, T_+) .*

Considering the word w in Fig. 1, we see that the carets numbered zero have type pairing (L_0, L_0) , which has weight 0. The carets numbered 1 have types (L_L, I_R) which contributes 2 to the weight of the word. The total weight of the word is easily computed

to be $0 + 2 + 4 + 2 + 2 + 1 + 1 + 4 + 3 + 1 + 0 = 20$. Thus, the length of w in the word metric $|w|_{\mathcal{F}}$ is 20.

The proofs in Section 4 rely heavily on this technique of Fordham. Namely, we use the fact that we can apply a generator to a given word, whose length we know, and the change in caret types, which is easily seen, exactly determines the change in word length.

2.2. Action of the generators on an element of F

We begin with a lemma from Fordham [8] which states under fairly broad conditions, that when applying a generator to a tree pair (T_-, T_+) exactly one pair of caret types will change. In Section 3, we construct a special family of elements which will provide the counterexamples to almost convexity for the standard two-generator generating set for F . These elements are constructed to satisfy the conditions of the lemma below.

Lemma 2.3 (Fordham [8, Lemma 2.3.1]). *Let (T_-, T_+) be a reduced pair of trees, each having $m + 1$ carets, representing an element $x \in \mathcal{F}$, and γ any generator of \mathcal{F} .*

- (1) *If $\gamma = x_0$, we require that the left subtree of the root of T_- is nonempty.*
- (2) *If $\gamma = x_0^{-1}$, we require that the right subtree of the root of T_- is nonempty.*
- (3) *If $\gamma = x_1$, we require that the left subtree of the right child of the root of T_- is nonempty.*
- (4) *If $\gamma = x_1^{-1}$, we require that the right subtree of the right child of the root of T_- is nonempty.*

If the reduced tree pair diagram for x also has $m + 1$ carets, then there is exactly one i with $0 \leq i \leq m$ so that the pair of caret types of caret i changes when γ is applied to x .

We now begin to understand geometrically the action of a generator of \mathcal{F} on a reduced tree pair (T_-, T_+) , and the corresponding change in normal form. We will generally assume that the conditions of Lemma 2.3 are met by the generic elements with which we begin.

Let C_R denote the caret which is the right child of the root caret R of T_- , and C_{RR} and C_{RL} the right and left carets, respectively, of C_R . Similarly, let C_L denote the left child of the root caret of T_- , and C_{LL} and C_{LR} its left and right children. Figures 2, 3 and 4 will be useful in understanding the geometric interpretation of the action of the generators on an element of F . In all of these figures, the letters a, b and c represent (possibly empty) subtrees of the given tree.

We first understand the action of the generator x_0^{-1} on a tree pair (T_-, T_+) representing an element $w \in F$. Consider w written in normal form as $w = x_{i_1}^{r_1} \cdots x_{i_n}^{r_n} x_{j_m}^{s_m} \cdots x_{j_1}^{s_1}$. Then the element $w x_0^{-1}$ is still in normal form (unless we are in the degenerate case where $x = x_0^m$). Recall from Section 2 that the exponent of x_0^{-1} in the normal form is one less than the number of left edges of the tree T_- . Thus, increasing the exponent of x_0^{-1} by 1 adds a left edge to T_- .

The numbering of the leaves and carets after this new edge is added must remain the same, since the normal form (and hence the exponents of the leaves) changes in a single place. Thus, with the extra edge in T_- , C_R becomes the new root caret. The left subtree

of C_R , which contains carets with smaller numbers than C_R , must become the right subtree of the old root caret, which is now at position formerly occupied by C_L . The left caret C_L is moved down and to the left and remains a left caret, now in the position formerly occupied by C_{LL} and so on. This tree transformation is also called a *counterclockwise rotation* or *left rotation* based at the root. Figure 2 shows the negative trees T^- for the elements w and $w x_0^{-1}$ and illustrates a counterclockwise rotation based at the root.

When we consider the action of x_0 on $w = (T^-, T_+)$, we can assume, according to Lemma 2.3, that T^- has at least two left edges, equivalently, that the exponent of x_0^{-1} in the normal form of w is at least 1. Applying the generator x_0 cancels one x_0^{-1} in the normal form. This corresponds to the tree T^- losing a left edge, and thus the caret C_L becomes the root caret and the former root caret R moves to the position of C_R . The initial right subtree of C_L becomes the left subtree of R in order to preserve the numbering of the carets. This is a clockwise (or right) rotation based at the root of T^- and is illustrated in Fig. 3.

It is more difficult to visually understand the action of x_1 and x_1^{-1} on the pair (T^-, T_+) corresponding to w , as it is more difficult to see how these generators change the normal form. Using the terminology given above, the following lemmas show that the generators x_1 and x_1^{-1} perform counterclockwise and clockwise rotations around the node C_R .

We begin with a lemma relating the action of x_1^{-1} on (T^-, T_+) to the normal form of the corresponding element $w \in \mathcal{F}$.




Fig. 2. Rotation at the root induced by applying x_0^{-1} to T^- .




Fig. 3. Rotation at the root induced by applying x_0 to T^- .

Lemma 2.4 (The normal form of $w x_1^{-1}$). *Let $w \in F$ be represented by the tree pair (T^-, T_+) , and have normal form $x_1^{r_1} \cdots x_{i_n}^{r_n} x_{j_m}^{s_m} \cdots x_{j_1}^{s_1}$. Then $w x_1^{-1}$ has normal form*

$$x_{i_1}^{r_1} \cdots x_{i_n}^{r_n} x_{j_m}^{s_m} \cdots x_{j_{q+1}}^{s_{q+1}} x_{j_q}^{-1} x_{j_q}^{s_q} \cdots x_{j_1}^{s_1}, \tag{1}$$

where we might have $j_{q+1} = j_q + 1$. If the root of T^+ has right and left subtrees T_R and T_L , respectively, then j_{q+1} is smallest leaf number in T_R .

Proof. We consider the proof in two cases. In the first case, if $j_1 \neq 0$ then $j_1 = 1$ and the expression $x_1^{r_1} \cdots x_{i_n}^{r_n} x_{j_m}^{s_m} \cdots x_{j_1}^{s_1} x_1^{-1}$ is in normal form. In this case, T^+ has a single left edge on the left side of the tree, with leaf labelled 0, and the first left leaf of the first right subtree will be labelled 1.

In the second case we assume that $j_1 = 0$. Then the relators in \mathcal{P} imply that $j_l = 1 + s_1 + s_2 + \cdots + s_l$, where l is the first index satisfying $j_{l+1} \geq 1 + s_1 + s_2 + \cdots + s_l$. It remains to show that this is the label of the leftmost leaf of the first right subtree of T^+ .

Let T_L and T_R be the left and right subtrees of the root caret of T^+ . We consider the number of interior carets in T_L . If T_L is empty, then we are in the first case discussed above.

If T_L has no interior carets, but is not empty, then the number of left edges in T_L is n , for some n , and thus the last leaf number in T_L is n as well. So the first leaf number in T_R is $n + 1$. Given this form of T^+ , we see that the normal form of x must end with $x_{j_2}^{s_2} x_0^n$ where $j_2 \geq n + 1$. Thus, using the relators to put x_1^{-1} into its proper position in the normal form, we see that it becomes x_{1+n}^{-1} , agreeing with the statement of the lemma.

If T_L has a single interior caret, then the total number of left edges of T_L is $n + 1$, where n again represents the length of the left side of T_L . The interior caret also adds an additional leaf, and thus the highest numbered leaf of T_L is $n + 1$. We know that x_1^{-1} becomes x_{n+1}^{-1} when it is moved left past the x_0^n . However, $n + 1$ is now the highest numbered leaf in T_L . The extra left leaf added by the single interior caret corresponds to a letter in the normal form of x whose index is smaller than $n + 1$, thus when the x_{n+1}^{-1} is moved left past this letter, it becomes an x_{n+2}^{-1} . Since there are no other interior carets in T_L , the next possible index of a letter in the normal form of x is $n + 2$. Thus x_{n+2}^{-1} is now in place in the normal form, so $j_2 = n + 2$, and $n + 2$ is the first leaf number of T_R , as required.

If T_L has two interior carets, then there are $n + 2$ left edges in T_L and the highest leaf number in T_L is $n + 2$. Moving left past x_0^n , the x_1^{-1} first becomes x_{n+1}^{-1} , as in the previous case. Again, we see that $n + 1$ is a leaf number in T_L . Then, since there is a single leaf numbered higher than $n + 1$, there are not enough leaves to have the remaining two carets have leaves numbered higher than $n + 1$. So the first interior caret must have a leaf with a lower number than $n + 1$, corresponding to a letter in the normal form of x with index smaller than $n + 1$. Thus x_{n+1}^{-1} must be moved left past this element as well, making it x_{n+2}^{-1} . Now, $n + 2$ is the highest leaf number in T_L , so the second interior caret must again appear before leaf number $n + 2$; that is, it corresponds to a letter in the normal form of x with index smaller than $n + 2$. Moving the x_{n+2}^{-1} past left this letter, we get x_{n+3}^{-1} . Since there are no more interior carets in T_L , there are no other letters in the normal form with index less than $n + 3$, so we must have x_{n+3}^{-1} in its place in the normal form. Again, we see that $n + 3$ is the first leaf number in T_R .

In summary, each additional interior caret adds a letter to the normal form with smaller index than $n + 1$; thus the x_1^{-1} must be moved left past these letters to obtain the normal form. We can continue this method to apply to an arbitrary number of interior carets in T_L , proving the lemma.

Lemma 2.5 (The normal form of $w x_1$). *Let w satisfy the conditions of Lemma 2.3 and have normal form $x_1^{r_1} \cdots x_{i_n}^{r_n} x_{j_m}^{s_m} \cdots x_{j_1}^{s_1}$. Then $w x_1$ has normal form:*

$$x_1^{r_1} \cdots x_{i_n}^{r_n} x_{j_m}^{s_m} \cdots x_{j_l}^{(s_l - 1)} \cdots x_{j_1}^{s_1}, \tag{2}$$

for some index j_l , in which case j_l is the smallest leaf number in the right subtree of T .

Proof. As in the proof of Lemma 2.4, we use the relators of \mathcal{P} to move x_1 to the generator x_j where $j = 1 + s_1 + s_2 + \cdots + s_l$, and l is the first index satisfying the inequality $j_{l+1} \geq 1 + s_1 + s_2 + \cdots + s_l$, or $j = 1 + s_1 + s_2 + \cdots + s_m$. From the proof of Lemma 2.4 we again know that j is the number of the leftmost leaf of the first right subtree of T .

According to Lemma 2.3, the left subtree of C_L is nonempty, so there is a leaf labelled a with exponent at least 2, i.e., there is an index $j_k = j$ in the normal form of w . Thus the exponent of x_{j_k} decreases by 1 because the x_1^{-1} cancels one $x_{j_k}^{-1}$ letter giving the normal form (2).

Lemma 2.6 (The action of x_1^{-1} on T). *The generator x_1^{-1} when applied to an element w of F represented by a tree pair (T, T_+) which satisfies the conditions of Lemma 2.3 leaves T_+ unchanged, and affects T as follows: C_{RR} becomes the right child of the root caret, and C_R becomes the left child of C_{RR} . All other carets remain unchanged.*

Proof. Let j be the number of the leftmost leaf in the right subtree of the root of T . It follows from Lemma 2.5 that the exponent of x_j in the normal form of w is increased by 1; that is, the exponent $E(j)$ of the leaf j is increased by 1, which means there is one more left edge emanating from C_R in T and terminating at j . Since the numbering of the carets is preserved, because the normal form changes in a single letter, and begins at the far left of the right subtree of the root caret, we see that C_R is now an interior caret. To preserve the

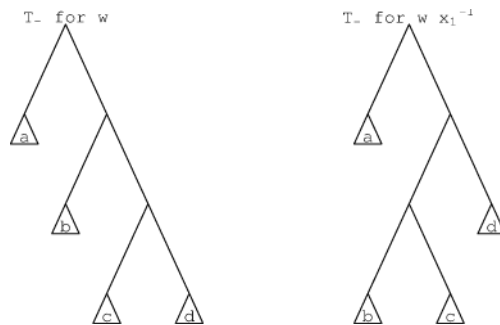


Fig. 4. Left rotation around C_R induced by applying x_1^{-1} .

numbering of the leaves and carets, the left subtree of C_{RR} must become the right subtree of C_R , because these carets are numbered higher than C_R but lower than C_{RR} . This leaves C_{RR} as the right child of the root caret. All remaining subtrees are left unchanged.

Lemma 2.7 (The action of x_1 on T^-). *The generator x_1 when applied to an element $w \in F$ represented by a tree pair (T^-, T_+) satisfying the conditions of Lemma 2.3 leaves T_+ unchanged, and in T^- , causes C_{RL} to become the right child of the root and C_R to become the right child of C_{RL} . All other carets remain unchanged.*

Proof. The normal form of $w x_1$ is of the form (2) given in Lemma 2.5. From Lemma 2.4 we know that the index j_l is the number of the leftmost leaf in the left subtree of C_R in T^- . From the change in normal form we see that the exponent of x_{j_l} decreases by 1 and thus in T^- the exponent $E(j_l)$ decreases by 1. Thus, there is one fewer left edge emanating from C_R ending in the leaf numbered j_l . Accordingly, the right subtree of C_{RL} is moved to the right side of T^- , without changing the numbering of the carets. Thus C_{RL} is now the right child of the root, and C_R is the left child of C_{RL} .

Notice that in all of the descriptions above, the tree T_+ is not affected by the action of a generator. This is not true in general for reduced tree pair diagrams not satisfying the conditions of Lemma 2.3. In general, T_+ can be affected in exactly three ways:

- (1) when T^- has a single left edge, and the generator is x_0 ,
- (2) when the left subtree of C_R of T^- is empty, and the generator is x_1 , or
- (3) if the generator is $x_i \in \{x_0^{\pm 1}, x_1^{\pm 1}\}$ and the pair of trees corresponding to x_i is not reduced.

We choose the family of words which will provide the counterexamples to almost convexity so that the conditions of Lemma 2.3 are always satisfied.

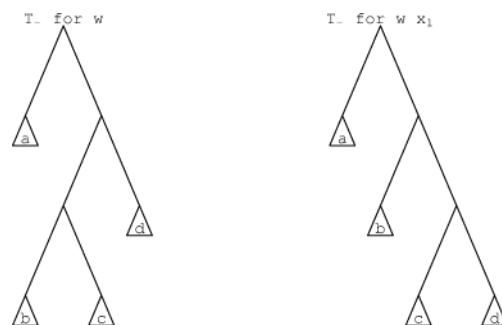


Fig. 5. Right rotation around C_R induced by applying x_1 .

3. A special family of elements

We define a family $\mathcal{F}(k)$, with integral $k \geq 2$, of elements of F which we will use to prove that F is not $AC(2)$, and thus not $AC(n)$. We first define what the negative tree T of an element $w \in \mathcal{F}(k)$ must be, and then define the positive tree T_+ so that w is given by the reduced tree pair (T, T_+) .

Let T_k be the balanced rooted binary tree with 2^k leaves; that is, the tree with every node on the first k levels having two children, as in Fig. 6.

For $w = (T, T_+)$ in the family $\mathcal{F}(k)$, we define T to be the tree T_{4k} . Note that this is a very bushy tree, and has at least $2k$ carets on the left side. Each of these left carets has a right subtree which is a complete tree with at least $k + 2$ levels. Similarly, T has at least $2k$ right carets, each of which has a left subtree which is a complete tree with at least $k + 2$ levels. There are a total of 2^{4k} leaves.

We construct the positive tree T_+ to have almost all carets of type L_L and R_{NI} , paired in a particular way with carets of T . Let $r = 2^{k-1} + 2^{k-2} - 1$ be the caret number of the first caret on the right side of T . Now let the tree T_+ correspond to the word $x_0^{r-2} x_1 x_s$, where s is $2^{4k} - 3$. Then T_+ will have 2^{4k} leaves, the same number as in T .

We now check that with these definitions, (T, T_+) forms a reduced tree pair diagram. As pictured in Fig. 7, there are only two carets in T_+ with two leaves: one with leaves numbered 1 and 2, and the other with leaves numbered $s = 2^{4k} - 3$ and $s + 1 = 2^{4k} - 2$. In T , it is easy to see that because it is a complete tree, caret number 0 has leaves numbered 0 and 1. Also in T , the highest numbered caret has leaves numbered $s + 1$ and $s + 2$. Thus, no reduction of carets occurs, and (T, T_+) is a reduced tree pair diagram.

In Section 2.2 above, the action of the generators of F on a generic element is discussed. We now describe the action of a generator on an element w of $\mathcal{F}(k)$, and more generally, the action of a sequence of generators on w . Let η be a word in the generators of F which has length strictly less than k , and let $w = (T, T_+) \in \mathcal{F}(k)$. We want to make sure that $w\eta$ still satisfies the conditions of Lemma 2.3. Because η is not longer than k , it can only affect a limited number of carets near the root of T . For example, if η is a power of x_1 , then each application of x_1 will rotate at the right child C_R of the root. The left subtree of C_R is, by construction, a complete tree with at least $k + 2$ levels. Thus, after performing k clockwise




Fig. 6. The balanced tree T_4 .




Fig. 7. Positive tree for a word $w \in (k)$.

rotations at the right child of the root, the resulting tree still satisfies the conditions of Lemma 2.3.

More generally, no matter what the sequence of generators in η is, the composition of rotations that η performs on T affects carets only within distance k of the root. Because of the fullness of the subtrees near the root of T , the resulting tree will still have carets in the appropriate locations to satisfy the conditions of Lemma 2.3. Because the exposed carets in T_+ are so far away from the root, we know that no reductions can happen during the course of applying η to w . Thus, Lemma 2.3 guarantees that only one caret is affected by each application of a generator.

In the following chart we summarize the possible change in word length when a generator of \mathcal{F} acts on an element $w\eta$ with $|\eta| < k$ and $w \in (k)$. The positive tree T_+ has been chosen carefully so that a caret in $w\eta$ affected by a generator is paired with one of only two possible types of carets in T_+ , an L_L or an R_{NI} .

| Generator | Original caret type | New caret type | Change in word length when paired with L_L | Change in word length when paired with R_{NI} |
|------------|---------------------|----------------|----------------------------------------------|-------------------------------------------------|
| x_0 | L_L | R_I | 1 | 1 |
| x_0^{-1} | R_I | L_L | 1 | 1 |
| x_1 | I_R | R_I | 1 | 1 |
| x_1^{-1} | R_I | I_R | 1 | 1 |

We see immediately from this chart that x_0 and x_0^{-1} will reduce the word length of $w \in (k)$ because of the caret pairings in w . It is also true from the chart that x_1 will reduce the length of the original word w . The two elements we will consider to contradict almost convexity will be wx_0 and wx_0^{-1} for $w \in (k)$. If the length $|w| = n + 1$, then the length of wx_0^{-1} and wx_0 will each be n . Furthermore, those two elements are distance 2 apart since there is an obvious path from wx_0 to w to wx_0^{-1} of length 2. That path, however,

does not lie in the ball of radius n . In the proof of Theorem 1.1, we will show that there is no short path from wx_0 to wx_0^{-1} which lies in the ball of radius n .

4. Almost convexity and

We now prove that F does not satisfy Cannon's $AC(2)$ condition, and obtain as a corollary that F does not satisfy $AC(n)$ for any integral $n \geq 2$.

The idea of the proof of Theorem 1.1 is the following. Assuming F satisfies the $AC(2)$ condition, we would obtain a constant k so that any two points in $B(n)$ at distance 2 from each other would be connected by a path of length at most k which remains in $B(n)$. Using this constant k , consider a point $w = (T^-, T_+) \in (k+2)$. The points wx_0 and wx_0^{-1} are both in $B(n)$ for $n = |w| - 1$ and are distance two apart. Thus, there would be a path γ of length at most k connecting them. We assume this path is oriented to go from wx_0 to wx_0^{-1} and we follow the position of the root caret R of T^- as it moves under the letters in the path γ . We know that in wx_0 the caret R has moved to the right side of the new negative tree. The main lemma to the proof of this theorem says that if at any time along the path γ the caret R becomes a left or an interior caret, then the path γ leaves $B(n)$ at that point.

Let $\gamma' = x_0\gamma x_0$ denote the loop based at w . The contradiction to almost convexity arises from the following: Since the word wx_0 has R as the right child of the root, and the word wx_0^{-1} has R as the left child of the root, the final x_0 in the path γ' would return R to the root position from the left. Thus, at some point along γ , the caret R would have changed from a right caret to a left or interior caret and at that point, the path γ would have left the ball $B(n)$.

We begin with the proof of the necessary lemma.

Lemma 4.1. *Let $w = (T^-, T_+) \in (k)$ with $|w| = n + 1$, and $\gamma' = x_0^m \gamma'' x_0$ be a loop based at w of length at most k , with m maximal. Let R be the root caret of T^- , and η the shortest prefix of γ'' so that in $wx_0^m \eta$ the caret R is not a right caret. Then the element $wx_0^m \eta$ is not in $B(n)$.*

Proof. First, note that the negative tree of the element wx_0^m has exactly m right carets which are paired with L_L carets, and we can number them as we move away from the root as $c_1, c_2, \dots, c_m = R$, with $c_1 < c_2 < \dots < c_m$. Since the numbering of the carets does not change when generators are applied, at the first point where R is not a right caret, then neither are any of the carets c_i .

In the statement of the lemma, we are not distinguishing between R becoming a left caret and R becoming an interior caret. This will not matter either for this proof or for the proof of Theorem 1.1 below.

The idea of the proof is to follow the path of each caret c_i as it is affected by different letters in the word η , and note the net change in word length. Note that when we apply a generator of F to a word of the form $w\chi$, where χ is a word in the generators of F of length at most k , only a single caret in the negative tree of $w\chi$ is affected. In general, there are times when this action can also affect a caret in the positive tree, but we have chosen

the form of elements of (k) carefully so that this is not the case, when applying strings of generators of length less than k .

Each caret c_i is originally paired with an L_L caret in the positive tree by construction, and since the positive tree will be unchanged, the positive part of these pairing types will not change. Consider all the letters in η which change the caret type of c_i . The last of these letters is either an x_0^{-1} changing c_i from a right caret to a left caret, or an x_1^{-1} changing c_i from a right caret to an interior caret. According to the chart in Section 3, this is a net change in word length of $+1$.

There are other letters in η which can affect the caret c_i . However, they must come in pairs, each pair leaving c_i as a right caret so that the final letter in η which affects it can change it to a left or interior caret. These pairs can be in one of two forms:

- (1) an x_0^{-1} which makes c_i a left caret followed later in η by an x_0 making it again a right caret, or
- (2) an x_1^{-1} making c_i an interior caret and an x_1 later in η making it again a right caret.

In either case, c_i is always paired with an L_L caret, and we see from the chart in Section 3 that the net change to the total word length corresponding to either of these pairs is always 0. Thus, as we consider the letters of η which change the caret type of all m of the c_i 's, we see that they contribute a total of $+m$ to the overall change in word length.

There may be letters in η which affect the types of carets other than the c_i . Suppose caret $d \neq c_i$ is a caret affected by a letter in η . We claim that we must have $d < c_i$ for some i , and thus d is also paired with an L_L caret. If $d > c_i$ for all i , then d would be a caret which appears after R . In order for η to affect a caret after R , the caret R would have had to have already moved from a right caret to a left or interior caret, contradicting our assumption about η . Thus, we have established the claim that $d < c_i$ for some i .

Given the initial form of $w \in (k)$, we see that d may begin as an interior caret, and be initially moved to a right caret by an element x_1 . From the chart in Section 3 we see that this changes word length by -1 . Since $d < c_i$ for at least one value of i , and all the c_i must be changed from right carets to non-right carets by the end of the path η , we must also have d changed from a right caret to a non-right caret. Thus the last letter in η affecting d is either an x_0^{-1} which changes d to a left caret or an x_1^{-1} which changes d back to an interior caret. From the chart in Section 3 we see that in either case, the change to the word length is $+1$ making the total contribution of these two letters in η zero.

There may be other letters in η which affect the caret d . They must form the same pairs as listed above of "intermediate" letters which can affect the c_i , and thus contribute a total word length change of zero.

The only other possibility for d is that it begins as a left caret, paired with an L_L caret for the same reasons as above. Then the initial letter in η affecting d must be an x_0 , making it a right caret. The final letter in η affecting d again is either an x_0^{-1} or an x_1^{-1} . Again, we see from the chart in Section 3 that the net change in word length coming from these two elements is 0. There can also again be intermediate pairs of elements affecting d of the same forms as given above, which also contribute 0 to the net change in word length.

Since every letter of η affects a single caret of w , each letter of η is one of the types listed above. So the total change in word length from wx_0 to $wx_0^m\eta$ is m . Given

the initial form of w , it is easy to see from the chart that $|wx_0^m| = |w| - m$. Thus $|wx_0^m \eta| = |w| - m + m = |w| = n + 1$ and $wx_0^m \eta$ is not in $B(n)$.

We are now ready to prove Theorem 1.1 using Lemma 4.1.

Proof of Theorem 1.1. Assume that F satisfies the $AC(2)$ condition with respect to the generating set $\{x_0, x_1\}$. Then there would be a constant k so that for every two points $x, y \in B(n)$ with $d(x, y) = 2$, there would be a path between them of length at most k lying completely inside $B(n)$.

Consider a point $w = (T^-, T_+) \in (k + 2)$ with $|w| = n + 1$. By construction, $|wx_0| = |wx_0^{-1}| = n$ and $d(wx_0, wx_0^{-1}) = 2$. The assumption of almost convexity guarantees a path γ from wx_0 to wx_0^{-1} which remains inside $B(n)$ whose length is bounded by k . Let $\gamma' = x_0 \gamma x_0$ be the loop based at w containing the path γ .

Let R be the root caret in T^- . The word wx_0 has R as the right child of the root, so the initial x_0 in the path γ' moves R to a right caret. The word wx_0^{-1} has R as the left child of the root, so the final x_0 in the path γ' must return R to the root position from the left. Thus, at some point along the path γ , the caret R must change from being a right caret to a left caret. So there is a minimal prefix η of γ so that in $wx_0 \eta$, the caret R is not a right caret. It then follows from Lemma 4.1 that $wx_0 \eta$ is not in $B(n)$, contradicting the assumption that F is $AC(2)$.

We immediately obtain the proof of Corollary 1.2.

References

- [1] M.G. Brin, C.C. Squier, Groups of piecewise linear homeomorphisms of the real line, *Invent. Math.* 44 (1985) 485–498.
- [2] K.S. Brown, R. Geoghegan, An infinite-dimensional torsion-free FP_∞ group, *Invent. Math.* 77 (1984) 367–381.
- [3] J. Burillo, S. Cleary, M.I. Stein, Metrics and embeddings of generalizations of Thompson’s group F , *Trans. Amer. Math. Soc.* 353 (4) (2001) 1677–1689 (Electronic).
- [4] J. Burillo, Quasi-isometrically embedded subgroups of Thompson’s group F , *J. Algebra* 212 (1) (1999) 65–78.
- [5] J.W. Cannon, W.J. Floyd, M.A. Grayson, W.P. Thurston, Solvgroups are not almost convex, *Geom. Dedicata* 31 (3) (1989) 291–300.
- [6] J.W. Cannon, W.J. Floyd, W.R. Parry, Introductory notes on Richard Thompson’s groups, *Enseign. Math.* (2) 42 (1996) 215–256.
- [7] J.W. Cannon, Almost convex groups, *Geom. Dedicata* 22 (2) (1987) 197–210.
- [8] B. Fordham, Minimal Length Elements of Thompson’s group F , PhD thesis, Brigham Young Univ., 1995.
- [9] P. Freyd, Splitting homotopy idempotents, in: *Proc. Conf. Categorical Algebra*, Springer, Berlin, 1966, pp. 173–176.
- [10] P. Freyd, A. Heller, Splitting homotopy idempotents ii, *J. Pure Appl. Algebra* 89 (1–2) (1993) 93–106.
- [11] V. Guba, M. Sapir, Diagram groups, *Mem. Amer. Math. Soc.* 130 (620) (1997) viii+117.
- [12] C.F. Miller, III, M. Shapiro, Solvable Baumslag–Solitar groups are not almost convex, *Geom. Dedicata* 72 (2) (1998) 123–127.
- [13] M. Shapiro, M. Stein, Almost convex groups and the eight geometries, *Geom. Dedicata* 55 (2) (1995) 125–140.

- [14] C. Thiel, Zur fast-Konvexität einiger nilpotenter Gruppen, Universität Bonn Mathematisches Institut, Bonn, 1992, Dissertation, Rheinische Friedrich-Wilhelms-Universität Bonn, Bonn, 1991.
- [15] R.J. Thompson, R. McKenzie, An elementary construction of unsolvable word problems in group theory, in: W.W. Boone, F.B. Cannonito, R.C. Lyndon (Eds.), Word Problems, Conference at University of California, Irvine, 1969, North-Holland, Amsterdam, 1973.