BOUNDING RIGHT-ARM ROTATION DISTANCES

SEAN CLEARY AND JENNIFER TABACK

Abstract. Rotation distance quantifies the difference in shape between two rooted binary trees of the same size by counting the minimum number of elementary changes needed to transform one tree to the other. We describe several types of rotation distance, and provide upper bounds on distances between trees with a fixed number of nodes with respect to each type. These bounds are obtained by relating each restricted rotation distance to the word length of elements of Thompson's group F with respect to different generating sets, including both finite and infinite generating sets.

1. Introduction

Rotation distance quantifies the difference in shape between two rooted binary trees of the same size by counting the minimum number of elementary changes needed to transform one tree to the other. Search algorithms are most efficient when searching balanced trees, which have few levels relative to the number of nodes in the tree. Thus one is often interested in calculating, or at least bounding, the number of these changes necessary to alter a given tree into one with a more desirable shape, such a balanced tree.

If we allow these elementary changes, called rotations, to take place at any node, we obtain ordinary rotation distance. This was analyzed by Sleator, Tarjan and Thurston [14], who proved that there is an upper bound of $2n-6$ rotations needed to transform one rooted binary tree with n nodes into any other. Furthermore, they showed that the $2n-6$ bound is achieved for large values of n and thus is the best possible upper bound. No efficient algorithm is known to compute rotation distance exactly, though there are polynomial-time algorithms of Pallo [11] and Rogers [13] which estimate rotation distance efficiently.

The first author acknowledges support from PSC-CUNY grant #66490, NSF grant DMS-0305545 and the hospitality of the Centre de Recerca Matemàtica.

The second author acknowledges support from NSF grants DMS-0305545 and DMS-0437481 and the hospitality of the Centre de Recerca Matemàtica.

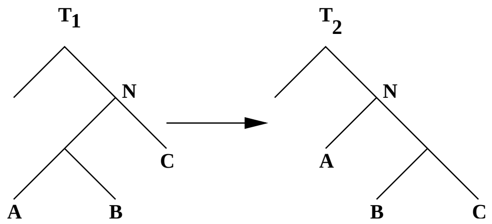

FIGURE 1. Right rotation at node N transforms T_1 to T_2; similarly, left rotation at node N transforms T_2 to T_1. The labels A, B and C represent (possibly empty) subtrees of T_1 and T_2.

Here we expand on the study of restricted rotation distance begun in [4] and [5]. Restricted rotation distance allows rotations only at the root node and the right child of the root node. Restricted rotation distance is related to the word length of elements of Thompson's group F with respect to its standard finite generating set. This is illustrated in [4, 5] and involves the interpretation of elements of F as pairs of finite binary rooted trees and Fordham's method [9] for computing the word length of an element of F with respect to that standard finite generating set directly from such trees. These methods not only give an effective algorithm to compute restricted rotation distance, but they also give an effective algorithm to find the appropriate rotations which realize this distance.

Right and left rotations at a node N of a rooted binary tree T are defined to be the permutations of the subtrees of T described in Figure 1. Right rotation at a node N transforms the original tree T_1, given on the left side of Figure 1, to the tree T_2 on the right side of Figure 1. Left rotation at a node is the inverse operation. In all that follows, T_1 and T_2 denote trees with the same number of leaves.

In this paper, we discuss generalizations and variations of restricted rotation distance, in which rotations are again only allowed at specified nodes of the tree. We relate these distances to distinct word metrics on Thompson's group F. We use this interpretation to exhibit asymptotically sharp linear bounds on the number of allowable elementary rotations needed to transform one tree with n nodes into another. These alternate definitions all allow rotations at the root node and at nodes connected to the root node by a path consisting entirely of right edges, that is, nodes that lie on the *right side* or *right arm* of the tree. The root node is considered to lie on the right side of the tree. While the original restricted rotation distance is well defined between any two trees with the same number of nodes, this is no longer necessarily the case when we allow rotations at other collections of nodes along the right side of the tree. Some transformations between trees cannot be accomplished with a specified set of rotations unless additional nodes are added to the trees– which is not permitted for rotation distance. Below, we describe when such a transformation is possible, and describe all pairs of trees for which rotation distance is defined.

The sharp upper bound on the restricted rotation distance between two trees, each with n nodes, obtained in [5] is $4n - 8$. Below, we consider allowing additional rotations along the right side of the tree and note that allowing rotations at any finite collection of nodes on the right side of the tree does not change the multiplicative constant of 4 in the upper bound. It is only when we allow an infinite set of rotations along the right arm of the tree that we obtain the multiplicative constant of 2 in the upper bound, analogous to ordinary rotation distance. These rotation distances and bounds are summarized in Table 1, where n is the number of nodes in each tree.

Type of distance:	**Rotations allowed at:**	**Symbol:**	**Upper Bound:**
Rotation distance	all nodes	d_R	$2n - 6$
Restricted rotation distance	root node and right child of the root node	d_{RR}	$4n - 8$
Restricted right arm rotation distance	root node and a finite collection $\mathcal{S}$ of nodes on the right side of the tree	$d^{\mathcal{S}}_{RRA}$	$4n - C$ some C
Right arm rotation distance	all nodes on the right side of the tree	d_{RA}	$2n - 2$
Restricted spinal rotation distance	root node and a finite collection $\mathcal{S}$ of nodes on both sides of the tree	$d^{\mathcal{S}}_{RS}$	$4n - C$ some C

TABLE 1. Summary of rotation distances between trees with n nodes.

Culik and Wood [8] allowed rotations only at nodes along the right side of the tree, and showed that under these conditions a sharp upper bound on this *right-arm rotation distance* between two trees with n nodes is $2n - 2$. Since the ordinary rotation distance between two such trees can be as much as $2n - 6$, it is remarkable that restricting rotations to the right side of the tree adds only four rotations to the upper bound. Below, we describe this in terms of the metric on Thompson's group F, and explain how allowing rotations at all nodes along the right arm of a tree corresponds to the word metric with respect to the standard infinite generating set for F.

Pallo studied right-arm rotation distance in [12], allowing rotations at all nodes along the right side of the tree. He described a process for computing right-arm rotation distance which we show below is equivalent to finding the word length in Thompson's group F with respect to the standard infinite generating set.

Each *node* in the trees we consider consists of a vertex and two downward directed edges. We will only consider finite, rooted binary trees with n nodes. These trees are called *extended binary trees* in Knuth [10] or *0-2 trees.* A vertex of valence one in a tree is called an *exposed leaf* or simply a *leaf.* The exposed leaves in a tree are numbered from left to right, beginning with zero. Our trees will have n nodes, which yield $n+1$ exposed leaves, numbered from 0 to n. A node with both of its leaves exposed is called an *exposed node*, and its leaves are termed *siblings* and those leaves are said to form a *sibling pair.* A node N which is attached to the right (respectively left) edge of a node M is called the *right (respectively left) child* of M. A node which has one edge on the left side of the tree is called a *left node.* A node which has one edge on the right side of the tree and is not the root node is called a *right node.* Nodes which are neither right nor left are called *interior nodes.* The union of left and right nodes in a tree is called the *spine* of the tree. A tree consisting of only the root node and $n-1$ right nodes is called the *all-right* tree with n nodes. An *ancestor* of a node is any node which lies along the shortest path between it and the root node.

The connection between Thompson's group F and restricted rotation distance is described below. Thompson's group F is combinatorially studied in two ways: via a finite presentation and an infinite presentation. Computing restricted rotation distance between two trees is related to computing the word length of the element of F described by those trees with respect to the standard finite generating set for the group F. Analogously, right-arm rotation distance corresponds to computing the word length of the element with respect to the word metric induced by the standard infinite generating set for F. Restricted right-arm rotation distances and restricted spinal rotation distances, defined below, relate to the word metric on F with respect to other finite generating sets for F.

2. Thompson's Group F

The connection between Thompson's group F and rotations at nodes of trees is described in [4] and [5], using the work of Fordham [9]. Here, we briefly describe this connection, and refer the reader to Cannon, Floyd and

Parry [3] for a survey of the properties of Thompson's group F, and the further connection between elements of F and pairs of binary rooted trees.

2.1. **The infinite presentation of Thompson's group F.** Thompson's group F has a presentation with an infinite number of generators and relations:

$$\mathcal{P} = \left\langle x_0, x_1, \dots | x_i^{-1} x_n x_i = x_{n+1}, \forall i < n \right\rangle.$$

In this presentation, there are normal forms for elements given by

$$x_{i_1}^{r_1} x_{i_2}^{r_2} \dots x_{i_k}^{r_k} x_{j_l}^{-s_l} \dots x_{j_2}^{-s_2} x_{j_1}^{-s_1}$$

with $r_i, s_i > 0$, where the indices satisfy $0 \leq i_1 < i_2 \dots < i_k$ and $0 \leq j_1 < j_2 \dots < j_l$. This normal form is unique for a given element if we further require the reduction condition 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]. The relations provide a quick and efficient method for rewriting words into normal form, and form a complete rewriting system, as described by Brown [1]. There is a natural shift homomorphism $\phi : F \to F$ where $\phi(x_i) = x_{i+1}$ which respects the relators, and the reduction from normal form to unique normal form is accomplished with a sequence of operations replacing words of the form $ux_i\phi(v)x_i^{-1}w$ with uvw, where $\phi(v)$ is a subword which contains only generators of index $i+2$ and higher.

We note that F can be generated by just x_0 and x_1 in the above presentation; the relators show that x_0 conjugates x_1 to x_2. Similarly, all higher-index generators are conjugates of x_1 by higher powers of x_0, as $x_n = x_0^{-(n-1)} x_1 x_0^{n-1}$. This leads to a finite presentation for F with generating set $\{x_0, x_1\}$. In fact, x_0 and any higher index generator are sufficient to generate the group. Two generators x_i, x_j with $i \neq j$ will generate an subgroup of F which is isomorphic to the entire group but which is the entire group only when one of i or j is 0.

We begin by proving that in the word metric arising from this infinite generating set, the normal form expressions are geodesic representatives for elements of F.

Lemma 2.1. *Let $w = x_{i_1}^{r_1} x_{i_2}^{r_2} \dots x_{i_k}^{r_k} x_{j_l}^{-s_l} \dots x_{j_2}^{-s_2} x_{j_1}^{-s_1}$ be in unique normal form, as described above. Then $x_{i_1}^{r_1} x_{i_2}^{r_2} \dots x_{i_k}^{r_k} x_{j_l}^{-s_l} \dots x_{j_2}^{-s_2} x_{j_1}^{-s_1}$ is a geodesic in the word metric arising from the standard infinite generating set $\{x_i\}$ of F.*

Proof. Suppose that $x_{i_1}^{r_1} x_{i_2}^{r_2} \dots x_{i_k}^{r_k} x_{j_l}^{-s_l} \dots x_{j_2}^{-s_2} x_{j_1}^{-s_1}$ was not a geodesic representative in this word metric. Then there is a shorter expression, not necessarily in normal form, which we call α, representing w in this infinite generating set. It is clear from the relations of $\mathcal{P}$ that the conversion of α into unique normal form can only preserve or decrease the length of α.

Thus, after converting α into normal form we have obtained a second expression of w in unique normal form shorter than the initial unique normal form for w, and we have a contradiction. □

2.2. **Tree pair diagrams for elements of Thompson's group F.** The group F has a geometric description in terms of equivalence classes of tree pair diagrams. A *tree pair diagram* is a pair of finite rooted binary trees with the same number of leaves. We write $w = (T_1, T_2)$ to denote the two trees comprising a pair representing w. The equivalence between the geometric and algebraic interpretations of F is described in [3], and examples of this equivalence and its connection with rotations are given in [7].

Given two trees with the same number of nodes T_1 and T_2, the word in normal form associated to $w = (T_1, T_2)$ is found as follows. The leaves of each tree are numbered from left to right, beginning with zero. The *leaf exponent* of a leaf numbered k is the integral length of the longest path starting at leaf k consisting entirely of left edges which does not touch the right side of the tree. The tree pair diagram (T_1, T_2) has an associated normal form $x_0^{f_0} x_1^{f_1} \dots x_n^{f_n} x_n^{-e_n} \dots x_1^{-e_1} x_0^{-e_0}$ where e_i is the leaf exponent of leaf i in tree T_1 and f_i is the leaf exponent of leaf i in T_2. We refer to T_1 as the *negative tree* of the tree pair diagram and T_2 as the *positive tree* of the tree pair diagram. An example of a tree with leaf exponents computed is given in Figure 2.

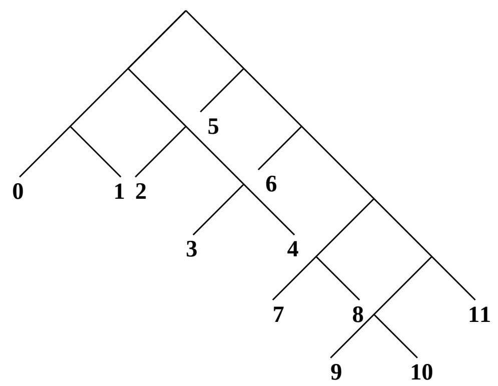

FIGURE 2. A tree whose leaves are numbered from left to right. The leaf exponents of the leaves, according to increasing leaf number, are 2,0,1,1,0,0,0,1,0,1,0, and 0.

An element of F is represented uniquely by a tree pair diagram satisfying the following reduction condition. A tree pair diagram (T_1, T_2) is *unreduced* if both T_1 and T_2 contain a node with two exposed leaves numbered i and $i+1$. A tree pair diagram which is not unreduced is *reduced.* Geometrically, any tree pair diagram has a unique reduced form that is obtained by successively deleting exposed nodes with identical leaf numbers from both trees, renumbering the leaves, and repeating this process until no further such reductions are possible. Elements of F are equivalence classes of tree pair diagrams, where the equivalence relation is that two tree pairs are equivalent if they have a common reduced form.

This tree pair reduction condition corresponds exactly to the combinatorial reduction condition given above to ensure uniqueness for words in normal form in the infinite presentation of F. That is, if leaves i and $i+1$ form a sibling pair in both T_1 and T_2, then in both cases, the leaf exponent of i will be non-zero and that for $i+1$ will be zero, as it is a right leaf. So the corresponding normal form will have both x_i and x_i^{-1} but no $x_{i+1}^{\pm 1}$, meaning that the normal form is not unique.

To perform the group operation on the level of tree pair diagrams, it may be necessary to use unreduced representatives of elements. Namely, to multiply (T_1, T_2) and (S_1, S_2), we create unreduced representatives (T_1', T_2') and (S_1', S_2') in which $T_2' = S_1'$, and write the product as the (possibly unreduced) element (T_1', S_2'). See [3] for examples of group multiplication using tree pair diagrams for elements of F.

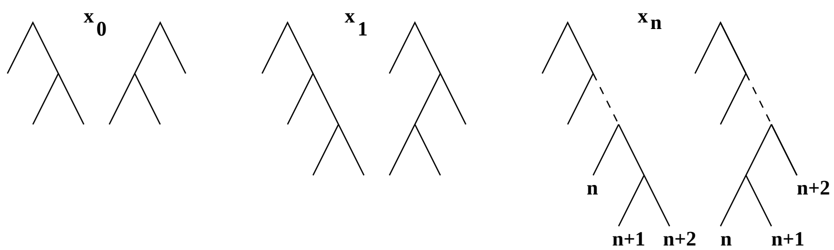

FIGURE 3. The tree pair diagrams corresponding to the generators x_0, x_1 and x_n of F.

The reduced tree pair diagrams associated to the generators x_0, x_1 and x_n are pictured in Figure 3. As explained in Lemmas 2.6 and 2.7 of [6], the generators x_0 and x_1 can be viewed in terms of rotations of rooted binary trees as well. The generator x_0 can be interpreted as a left rotation at the root of the left tree in the pair, yielding the right tree in the pair. Similarly, the generator x_1 performs a left rotation at the right child of the root node, transforming the left tree in the pair to the right one. The inverses x_0^{-1} and x_1^{-1} perform right rotations at the root node and right child of the root node, respectively.

One complication that may arise when using the geometry of the tree pair diagrams to understand rotation distance involves the possibility of adding nodes to a tree in order to perform group multiplication. Since elements of Thompson's group are equivalence classes of tree pair diagrams, we can always multiply any group element w by any group generator g. It is possible that we may have to add nodes to the reduced tree pair diagram for w in order to perform this multiplication. From the standpoint of group theory, the reduced and unreduced tree pair diagrams are interchangeable. When considering rotation distance, we are not allowed to change the number of carets in the starting tree. Thus certain rotations, corresponding to multiplication by specific generators, may not be permitted when calculating rotation distance.

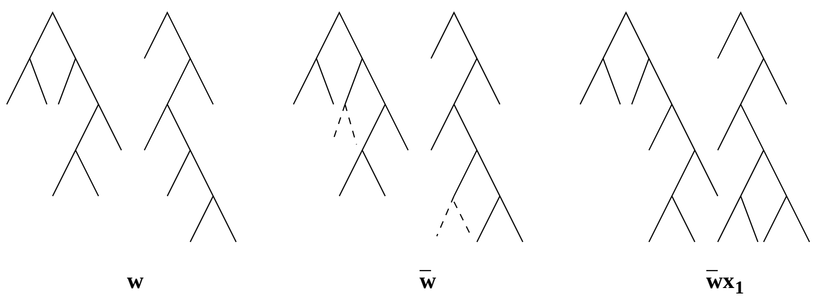

FIGURE 4. In order to multiply the tree pair diagram representing w by the generator x_1, we form the unreduced representative $\overline{w}$ of w in which the dashed node is added to both trees. Only then are we able to form the product wx_1.

For example, we cannot perform a right rotation at the right child of the root to either of the trees in the tree pair diagram for x_0 as shown in Figure 3 because in both of these trees, the left subtree of the right child of the root node does not exist. As an element of Thompson's group, we can enlarge any pair of trees to be able to multiply by any generator. A typical such application is shown in Figure 4 where a node is added to a tree to be able to perform the desired rotation. The tree pair diagram w does not have a left child of the right child of the root, so performing a right rotation at the right child of the root is not possible. However, the word w' which represents the same element of F does have a left child of the right child of the root and it is possible to perform the right rotation at the right child of the root there. We obtain w' by adding an additional node (indicated by dashing) to leaf number 3 in both trees of the tree pair diagram.

To describe when it is necessary to add a node to a tree to perform a particular rotation, we make the following definitions. We say *a right rotation at the root can be applied* to a tree T if the left subtree of the root of T is non-empty. Similarly, we say *a left rotation at the root can be applied* to a tree T if the right subtree of the root of T is non-empty and we also adopt this terminology when performing rotations at other nodes along the right side of the tree.

Understanding when rotations can be performed on trees helps us develop the connection between rotations of trees and right multiplication by generators of F. If, for example, we have a tree pair diagram (T_1, T_1) representing the identity and we can perform a left rotation at the root to T_1 to obtain x_0T_1, then the new tree pair diagram (T_1, x_0T_1) is the tree pair diagram representing the word x_0 in F, and similarly the new tree pair diagram (x_0T_1, T_1) is tree pair diagram representing the word x_0^{-1} in F.

We note that in F, multiplication by a generator may result in an unreduced tree pair diagram. So during the course of a long sequence of multiplications by generators of F, the number of nodes in the reduced tree pair diagram representing the partial products may fluctuate– rising when needing to add one or more nodes to apply a generator, and falling when multiplying by a generator results in an unreduced tree pair diagram. To understand rotation distance, however, as we apply a sequence of rotations to a single tree, we do not allow the number of nodes in the tree to change.

The link between restricted rotation distance and Thompson's group F is the word metric on F with respect to the generators $\{x_0, x_1\}$. Given two rooted binary trees T_1 and T_2 with the same number of nodes, we consider a minimal length word in $x_0^{\pm 1}$ and $x_1^{\pm 1}$ representing the element $w = (T_1, T_2) \in F$. As described in [4], this word gives a minimal sequence of rotations at the root and right child of the root which transform the tree T_1 into the tree T_2. It follows from Fordham [9] that these minimal words which transform one tree into the other maintain a constant number of carets at each stage in the sequence of rotations. The issue of certain rotations altering the number of nodes in the tree does not arise in the case of restricted rotation distance.

More precisely, suppose that $w \in F$ is given by the tree pair diagram (T_1, T_2), and a minimal length representative for w is $g_1g_2 \dots g_n$, where each $g_i \in \{x_0^{\pm 1}, x_1^{\pm 1}\}$. Then the tree pair diagram $(T_1, g_n \dots g_2g_1T_1)$ will represent w, and we can think of the sequence of generators $g_n \dots g_2g_1$ as a sequence of rotations which transforms T_1 to T_2. At each stage of this process, we will be able to perform the rotation corresponding to the generator g_{i+1} to the tree $g_i \dots g_2g_1T_1$ without adding additional nodes. There may be reductions possible to tree pair diagrams, or equivalently to the normal forms, during this process, but from the standpoint of rotation distance we do not want to take advantage of these reductions. Instead, we keep the number of nodes constant at each stage.

Equivalently, we can think of $(g_n \dots g_2 g_1 T_1, T_2)$ as a representative of the identity and witness the transformation of T_1 to T_2 by considering the sequence of tree pair diagrams

$$(T_1, T_2), (g_1 T_1, T_2), \dots, (g_n \dots g_2 g_1 T_1, T_2).$$

Below, we consider other possible locations for rotations to occur, and thus exploit the link to Thompson's group F, but now considering other generating sets for F. We assign a level to each node in the tree as follows. The root node is defined to have level zero. The level of a node N is the number of edges in a minimal length path connecting N to the root node. Writing the generators x_n for $n > 1$ via the relators $x_n = x_0^{-(n-1)} x_1 x_0^{n-1}$, we relate each generator to the following rotation of a tree T. We denote the all-right tree with the appropriate number of nodes by $*$. Group multiplication must be between a pair of elements, and each element corresponds to a pair of trees, so we use the tree $*$ as the positive tree corresponding to T. The product of the generator x_n and the tree pair diagram $(T, *)$ performs a right rotation to T at the node at level n along the right arm of T. In all that follows, when we describe a generator as inducing a rotation on a single tree T rather than on a tree pair diagram, we are forming the product with the pair $(T, *)$ as above.

3. Metrics on F and rotation distances

3.1. **Relation to the word metric.** In [5], the word length with respect to the finite generating set $\{x_0, x_1\}$ of F is used to compute the restricted rotation distance between a pair of trees, using techniques of Fordham [9]. Fordham developed a method for computing the exact length of an element of F directly from the reduced tree pair diagram representing that element.

Definition 3.1. *If T_1 and T_2 are trees with the same number of nodes, we define the* restricted rotation distance *$d_{RR}(T_1, T_2)$ as the minimal number of rotations required to transform T_1 to T_2, where rotations are allowed at the root and the right child of the root.*

Restricted rotation distance is well-defined for any two trees with the same number of leaves, as shown in [4]. We then obtain the following sharp bound on restricted rotation distance.

Theorem 3.2 ([5], Theorems 2 and 3)**.** *Given two rooted binary trees T_1 and T_2 each with n nodes, for $n \geq 3$, the restricted rotation distance between* *them satisfies* $d_{RR}(T_1, T_2) \leq 4n - 8$*. Furthermore, for all* $n \geq 3$*, there are trees* T_1' *and* T_2' *with* n *nodes realizing this bound, with* $d_{RR}(T_1', T_2') = 4n-8$*.*

Intermediate between the two-element generating set $\{x_0, x_1\}$ and the infinite generating set $\{x_0, x_1, \ldots\}$ are other finite generating sets of the form $\{x_0, x_{m_1}, \cdots, x_{m_k}\}$, where we arrange the indices of the generators in increasing order. Analyzing the infinite generating set corresponds to allowing all rotations along the right side of the tree. Finite generating sets correspond to allowing finite collections of rotations at the root node and other nodes along the right side of the tree.

Definition 3.3. *Let* $\mathcal{S} = \{x_0, x_{m_1}, \cdots, x_{m_l}\}$ *be a finite subset of the infinite generating set for* F *and* T_1 *and* T_2 *be trees with the same number of leaves. We define* $d^{\mathcal{S}}_{RRA}(T_1, T_2)$*, the* restricted right-arm rotation distance with respect to $\mathcal{S}$*, as the minimal number of rotations required to transform* T_1 *to* T_2*, where the rotations are only allowed at levels* $0, m_1, \ldots m_{l-1}$ *and* m_l *along the right side of the tree.*

We will see below that unlike restricted rotation distance, restricted right-arm rotation distance may not be defined between all pairs of trees with the same number of nodes. We use the notation $|\cdot|_{\mathcal{S}}$ to denote the word length of an element of F with respect to the generating set $\mathcal{S}$. We now relate the restricted right-arm rotation distance $d^{\mathcal{S}}_{RRA}(T_1, T_2)$ to $|(T_1, T_2)|_{\mathcal{S}}$.

Consider two trees T_1 and T_2 each with n nodes. The word length of the element $w = (T_1, T_2) \in F$ with respect to a generating set $\mathcal{S}$ is the length of the shortest expression for w in that generating set. However, when considering the corresponding rotations to the tree pair diagram for w, we have no analogue of Fordham's proof that a minimal length representative in these generators can be constructed while maintaining a constant number of nodes in each tree. Thus, it may be possible that a minimal length representative for $w = (T_1, T_2) \in F$ with respect to $\mathcal{S}$ includes some rotations which would require the addition of nodes to the trees and are thus not permitted. Therefore, we see that the word length $|(T_1, T_2)|_{\mathcal{S}}$ provides only a lower bound on the rotation distance $d^{\mathcal{S}}_{RRA}(T_1, T_2)$, when this rotation distance is defined. If this word length corresponds to a sequence of rotations in which the number of nodes remains constant at each intermediate step, then we have computed the actual restricted right-arm rotation distance between the two trees. These cases will be addressed below.

For example, we consider the trees shown in Figure 5. The desired transformation from the top left tree T_1 drawn in solid lines to the top right tree T_2 drawn in solid lines would be given by x_1, a single left rotation at the right child of the root. But if the permitted locations for rotation are only at the root (corresponding to the generator $x_0^{\pm 1}$) and the right child of the right child of the root (corresponding to the generator $x_2^{\pm 1}$), it will be impossible to accomplish the desired transformation without adding additional nodes, using any sequence of those allowed rotations and the corresponding restricted right-arm rotation distance is not defined between those two trees. If we are permitted to add a node to the leftmost leaf of each tree, as shown with the dashed nodes, to obtain the related problem of transforming the new tree T_1' into T_2' (drawn including the dashed nodes) then the transformation would be possible using only the allowed rotations. The unreduced form of the top tree pair diagram (T_1', T_2') drawn including the dashed node is $x_0 x_2 x_0^{-1}$ which reduces to x_1 in the usual manner, if desired. If rotations are permitted at the root and right child of the root, the rotations that transform T_1 to T_2 are exactly the same as those to perform the transformation from T_1' to T_2' and the added dashed node is simply carried along intact. However, if we are only permitted to rotate at the root and right child of the right child of the root, the added node is essential in allowing that transformation, though it does take two additional steps. We cannot transform T_1 to T_2 but we can easily transform T_1' to T_2' by rotating rightwards at the root, leftwards at the right child of the right child of the root, and then leftwards at the root, as pictured.

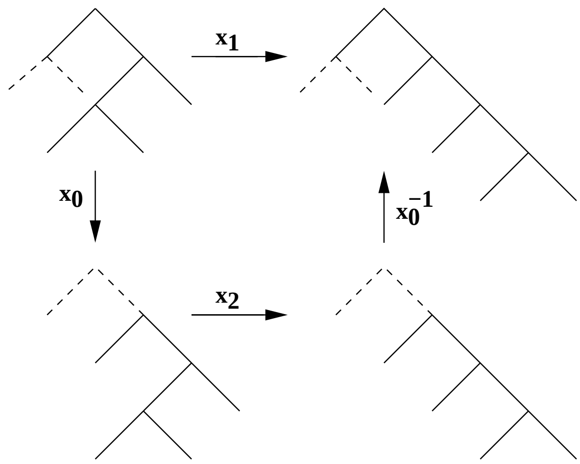

FIGURE 5. A right rotation at the right child of the root performed on the reduced tree by x_1 and on the partially reduced tree by $x_0 x_2 x_0^{-1}$.

We can describe exactly when a tree T_1 can be transformed into T_2 without adding nodes with respect to a specified set of allowed rotations along the right-arm of the tree, that is, when the restricted right arm rotation distance is defined. First, we consider the case when the word in normal form associated to (T_1, T_2) is already reduced; that is, when (T_1, T_2) is a reduced tree pair diagram.

Lemma 3.4. *Let $\mathcal{S} = \{x_0, x_{m_1}, \cdots, x_{m_l}\}$ be a generating set for F with $0 < m_1 < m_2 < \cdots < m_l$. We consider the corresponding restricted right-arm rotation distance $d^{\mathcal{S}}_{RRA}$, where rotations are allowed at nodes at levels $0, m_1 \ldots m_l$ on the right side of the tree. Suppose T_1 and T_2 are finite rooted binary trees with the same number of nodes forming a reduced tree pair diagram $w = (T_1, T_2) \in F$ with unique normal form given by*

$$x_{i_1}^{r_1} x_{i_2}^{r_2} \ldots x_{i_k}^{r_k} x_{j_l}^{-s_l} \ldots x_{j_2}^{-s_2} x_{j_1}^{-s_1}.$$

If $x_t^{\pm 1}$ for $1 \leq t \leq m_1 - 1$ appears in this normal form, then the restricted right-arm rotation distance $d^{\mathcal{S}}_{RRA}(T_1, T_2)$ is not defined. Conversely, if no $x_t^{\pm 1}$ for $1 \leq t \leq m_1 - 1$ appears in the unique normal form, then the restricted right-arm rotation distance $d^{\mathcal{S}}_{RRA}(T_1, T_2)$ is defined.

We note that if $m_1 = 1$ then we have included both of the generators used in restricted rotation distance, which together suffice to perform any transformation and then $d^{\mathcal{S}}_{RRA}$ will always be defined.

Proof. We recall that the leaf exponent of the leaf numbered n in a tree is the length of the maximal path of left edges from leaf n which does not reach the right side of the tree. Observe that the leaf exponent that changes as a result of a rotation at the node at level h on the right arm of the tree corresponds to the leftmost leaf in the left subtree of the node where the rotation occurs.

First, we suppose that $x_t^{\pm 1}$ for $1 \leq t \leq m_1 - 1$ appears in the unique normal form for (T_1, T_2). So t appears as a leaf label of a left leaf of a node in either in T_1 or T_2 or possibly both. If the restricted right-arm rotation distance $d^{\mathcal{S}}_{RRA}$ is defined, then the sequence of rotations transforming T_1 into T_2 does not change the number of nodes in the tree at any intermediate step and thus no leaves are added or removed during this process. We consider the leaf numbers whose exponents can be affected by rotations at the permitted nodes. Rotations are permitted at the root node and at levels m_i along the right side of the tree. Rotations at the root can affect only the exponent of leaf zero, as it will be the leftmost leaf in the left subtree attached at the root node. Other rotations can affect the exponents of leaves which are the leftmost leaves of left subtrees of right nodes at levels m_1 and lower. The left subtree of the right node at level h will have leaves numbered at least h, so if $t < m_1$, then no rotation at level m_i can affect the exponent of leaf t. If leaf t has different exponents in T_1 and T_2, since the allowed rotations cannot change its exponent, T_1 cannot be transformed into tree T_2 by the allowed rotations. If leaf t is present in both trees with the same exponent, then since w is in unique normal form, the exponent of leaf $t+1$ must also be non-zero in at least one of the trees. Moreover, leaves numbered t and $t+1$ belong to the same left subtree of a right node. Thus none of the allowed rotations can affect the leaf exponent of leaf $t+1$ as well. We iterate this argument with leaves $t+1$ and $t+2$. Thus, we see that if any $x_t^{\pm 1}$ with $1 \leq t \leq m_1 - 1$ appears, then the two trees cannot be connected by any sequence of the allowed rotations without the addition of extra nodes.

Conversely, if $x_t^{\pm 1}$ for $1 \leq t \leq m_1 - 1$ do not appear in the normal form, then we can rotate T_1 rightwards at the root by application of an appropriate power of x_0^k so that all of the nontrivial subtrees then hang from the right arm of the tree at heights m_1 and greater. We can then use x_0, x_{m_i} and conjugates of x_{m_i} by powers of x_0 to rotate the tree to an all-right tree, just as in the infinite generating set, without adding any additional nodes. So we can transform T_1 to the all-right tree, and then from the all-right tree, we can again use x_0, x_{m_1} and conjugates of x_{m_1} by powers of x_0 (and possibly other x_{m_i}, if desired) to transform the all-right tree to T_2 without adding additional nodes. Thus, $d^{\mathcal{S}}_{RRA}(T_1, T_2)$ is defined. There may be more efficient ways of accomplishing this transformation but it is clear that there is at least one way of doing it without adding additional nodes, so the restricted right-arm rotation distance is defined. □

To understand the case where (T_1, T_2) is an unreduced tree pair diagram, and thus we do not obtain the unique normal for the element directly from the leaf exponents, we introduce the notion of *partial reduction*. Partial reduction is similar to ordinary reduction except that we do not want to remove left nodes common to both trees. Stated algebraically, it means that if the normal form for the element contains instances of x_0 and x_0^{-1} but not $x_1^{\pm 1}$, we do not simplify the expression, as we do when x_k and x_k^{-1} appear but not $x_{k+1}^{\pm 1}$ for $k > 0$. The presence of these additional left nodes may allow us to perform rotations which would not be permitted otherwise without increasing the number of carets in the trees. This phenomenon occurs in the tree pairs shown in Figure 5.

Definition 3.5. *A word w in F in normal form is* partially reduced *if it is of the form $x_{i_1}^{r_1} x_{i_2}^{r_2} \dots x_{i_k}^{r_k} x_{j_l}^{-s_l} \dots x_{j_2}^{-s_2} x_{j_1}^{-s_1}$ with $0 \leq i_1 < \ldots < i_k$ and $0 \leq j_1 < \ldots < j_l$, with r_n and s_n all positive, and if we further require the partial reduction condition that for $i > 0$, when both x_i and x_i^{-1} occur, so does at least one of x_{i+1} or x_{i+1}^{-1}.*

For any word w in (not necessarily unique) normal form, there will be a maximal length word w' satisfying the partial reduction condition which we can easily obtain using the procedure described above.

The partial reduction allows us to prove the following lemma, which describes when one given tree can be transformed into another with respect to a specified set of rotations, when the initial tree pair diagram is unreduced. The proof is identical to that of Lemma 3.4.

Lemma 3.6. *Let $\mathcal{S} = \{x_0, x_{m_1}, \cdots, x_{m_l}\}$ be a generating set for F with $0 < m_1 < m_2 < \cdots < m_l$. We consider the corresponding restricted right-arm rotation distance $d^{\mathcal{S}}_{RRA}$ where rotations are allowed at the root node and at right nodes of levels $m_1 \dots m_l$. Suppose T_1 and T_2 are finite rooted binary trees with the same number of nodes forming a tree pair diagram $w = (T_1, T_2) \in F$ and that w has the partially reduced normal form of maximum length given by*

$$w' = x_{i_1}^{r_1} x_{i_2}^{r_2} \dots x_{i_k}^{r_k} x_{j_l}^{-s_l} \dots x_{j_2}^{-s_2} x_{j_1}^{-s_1}.$$

Then if $x_t^{\pm 1}$ for $1 \leq t \leq m_1 - 1$ appears in this partially reduced normal form, then the restricted right-arm rotation distance $d^{\mathcal{S}}_{RRA}(T_1, T_2)$ is not defined. Conversely, if no $x_t^{\pm 1}$ for $1 \leq t \leq m_1 - 1$ appears in this partially reduced normal form, then the restricted right-arm rotation distance $d^{\mathcal{S}}_{RRA}(T_1, T_2)$ is defined.

When rotations at all nodes along the right side of the tree are allowed, we obtain the right-arm rotation distance d_{RA} studied by Culik and Wood [8] and Pallo [12]. This situation is analogous to restricted rotation distance, which considers only the rotations corresponding to the generators x_0 and x_1, because the word length once again yields the exact rotation distance.

Proposition 3.7. *Let $\mathcal{I}$ denote the standard infinite generating set for F, and T_1 and T_2 be binary trees, each with n nodes. Then*

$$d_{RA}(T_1, T_2) = |(T_1, T_2)|_{\mathcal{I}}.$$

Proof. We will assume that the tree pair diagram (T_1, T_2) is reduced. If it is not, we form the tree pair diagram (T_1', T_2') representing the same group element which is reduced. The rotations necessary to transform T_1' into T_2' will also transform T_1 into T_2, since no additional rotations are necessary to alter the nodes which cause T_1 and T_2 to be unreduced. The nodes which were removed during the reduction are identical in both trees and are carried along unchanged during the rotations which transform T_1' to T_2'. The leaf exponent method of associating the unique normal form to the tree pair diagram described above shows that each tree provides one part of the normal form; in the pair (T_1, T_2) the tree T_1 corresponds to the terms with negative exponents and T_2 to those with positive exponents. We thus write the normal form as the product PN, where N contains the generators with negative exponents, and P those with positive exponents.

We see that N is a word which rotates the tree T_1 into the all-right tree without requiring the addition of any nodes, and the subword P is a string of generators which rotates the all-right tree into the tree T_2.

Thus we see that a lower bound for right arm rotation distance is $|(T_1, T_2)|_{\mathcal{I}}$, and an upper bound is given by combining the length of the strings P and N. It follows from Lemma 2.1 that $|(T_1, T_2)|_{\mathcal{I}} = |(T_1, *)|_{\mathcal{I}} + |(T_2, *)|_{\mathcal{I}}$ where $*$ is the all-right tree with n nodes, proving the proposition. □

3.2. **Bounds on restricted rotation distances.** Now that we have described the relationship between the different rotation distances and word lengths in F, we obtain numerical bounds on these rotation distances as summarized in Table 1. We note that word length of an element of F computed with respect to a generating set of the form $\mathcal{S}$ given above has the potential to be much shorter than the word length of the same element computed with respect to the generating set $\{x_0, x_1\}$. Thus we might expect significantly smaller asymptotic upper bounds on restricted right-arm rotation distance than on restricted rotation distance. In fact, this is not the case, and the difference between the upper bounds on the two rotation distances is at most a constant.

The goal of this section is to prove that the multiplicative constant of 4 in the upper bound on restricted right-arm rotation distance cannot be improved upon. That is, both the restricted rotation distance and the restricted right-arm rotation distance between two trees with n nodes each, when defined, are bounded above by $4n$ minus a constant. This constant depends upon the particular finite set of rotations permitted. These bounds are shown to be sharp for restricted rotation distance in [5]. We show below that they are asymptotically sharp for restricted right-arm rotation distance. While allowing additional rotations may shorten the restricted right-arm rotation distance between certain pairs of trees, asymptotically the worst-case scenario differs from restricted rotation distance by an additive constant. One way to improve the multiplicative constant of 4 is to allow rotation at an infinite collection of nodes along the right side of the tree, in which case the multiplicative constant may decrease to 2.

The necessity of the constant 4 is shown in two steps. We first show that the restricted right-arm rotation distance, when defined, is always bounded above by $4n$, where n is the number of nodes in either tree. We then show that there are words which realize this bound, up to an additive constant.

Proposition 3.8. *Let $\mathcal{S} = \{x_0, x_{m_1}, x_{m_2}, \dots x_{m_l}\}$ be a generating set for F with $0 < m_1 < m_2 < \cdots < m_l$, and let $d^{\mathcal{S}}_{RRA}$ the corresponding restricted right-arm rotation distance. Let T_1 and T_2 be binary trees, each with n nodes, for which $d^{\mathcal{S}}_{RRA}$ is defined. Then*

$$d^{\mathcal{S}}_{RRA}(T_1, T_2) \leq 4n.$$

Proof. The case where $m_1 = 1$ is already addressed by the analysis of ordinary rotation distance, described in [5]. We consider the element $w = (T_1, T_2) \in F$, where T_1 and T_2 are trees for which the relevant restricted right-arm rotation distance $d^{\mathcal{S}}_{RRA}$ is defined, and assume that $m_1 > 1$.

Case 1: The tree pair diagram (T_1, T_2) is reduced.

In this case, we know that the normal form of w contains no generators $x_t^{\pm 1}$ for $1 \leq t \leq m_1 - 1$. In addition, this normal form can contain x_0 or x_0^{-1} but not both. If both x_0 and x_0^{-1} were present in the normal form with no $x_1^{\pm 1}$ generator, then the normal form could be reduced. We can assume by symmetry that the normal form for w contains x_0^{-k} but no factors of x_0.

Using the correspondence between the normal form and the leaf exponents in the trees T_1 and T_2, we see that the leaves of both trees numbered from 1 through $m_1 - 1$ are either exposed right leaves of left nodes or exposed left leaves of right nodes. In T_1, denote the (possibly empty) subtrees of the left and right nodes by $A_1, A_2, \cdots, A_n$, where the smallest leaf number in A_1 is m_1. Similarly, in T_2 denote these subtrees by $B_1, B_2, \cdots, B_m$, where the smallest leaf number in B_1 is m_1.

Let $w' = w x_0^k$, so that the tree pair diagram (S_1, T_2) of w' has tree S_1 containing a single left node, namely the root node, and $m_1 - 1$ right nodes with exposed left leaves, followed by right nodes having $A_1, \cdots, A_n$ as their left subtrees. The pair (S_1, T_2) has the form given in Figure 6.

We consider the element $v \in F$ which has tree pair diagram (R_1, R_2), where R_1 has a single left node, namely the root node, and the left subtree of the right node at height i is A_i. The tree R_2 is defined analogously, using the subtrees B_i from the original tree T_2. Since restricted rotation distance is well defined for all trees with the same number of nodes, we apply Theorem 3.2 to obtain the bound $d_{RR}(R_1, R_2) \leq 4(n - (m_1 - 1)) - 8$. This restricted rotation distance is realized by a string α of the generators $\{x_0^{\pm 1}, x_1^{\pm 1}\}$.

We define a string of generators α' by replacing each instance of $x_1^{\pm 1}$ in α with $x_{m_1}^{\pm 1}$. Then this string of generators exactly produces the tree pair diagram (S_1, T_2). Since the number of nodes in each tree remains constant as each generator from α is applied to create (R_1, R_2), the same is true as we multiply the generators in α' to create $w' = (S_1, T_2)$.

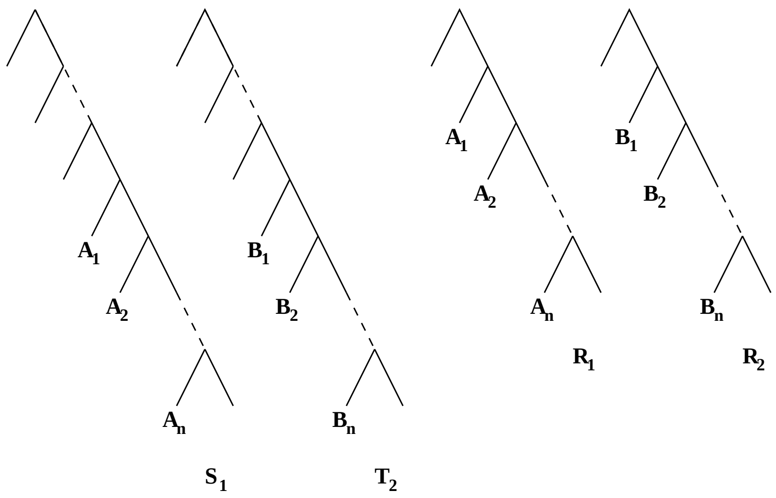

FIGURE 6. The tree pair diagrams for the words used to show that the restricted right-arm rotation distance is bounded above by $4n$.

Thus the restricted right-arm rotation distance with respect to $\mathcal{T} = \{x_0, x_{m_1}\}$ is bounded as follows:

$$d^{\mathcal{T}}_{RRA}(S_1, T_2) \leq 4(n - (m_1 - 1)) - 8.$$

Now we note that $w = w' x_0^k$, and since there were initially $k+1$ left nodes in the tree T_1, the number of nodes in each tree remains constant during these successive multiplications by x_0^{-1}. Thus the string $\alpha' x_0^{-k}$ realizes the restricted rotation distance between the trees T_1 and T_2.

If $k \leq m_1 - 1$, then the left nodes which are changed to right nodes under multiplication by x_0^{-k} do not appear in R_1 and R_2, and so are not represented in the upper bound given above. Thus, when the rotation distance is increased by k, we trivially extend the bound to

$$d^{\mathcal{T}}_{RRA}(T_1, T_2) \leq 4n - 8.$$

Since adding additional generators to the generating set, or equivalently allowing rotations at additional nodes, can only decrease the rotation distance, the upper bound still holds when we consider the entire generating set $\mathcal{S}$.

If $k \geq m_1$, then the left nodes which are changed to right nodes by this multiplication by x_0^{-1} are of two types: those with exposed left leaves numbered from 1 to $m_1 - 1$, and those with left subtrees of the form A_i.

The first type of right node is not counted in the upper bound given above, and thus we increase the number of nodes in the bound by $m_1 - 1$ to (more than) account for the additional generators.

The right nodes of the second type, with left subtrees of the form A_i, are already counted in the bound given above. However, we recall that the word α' which realizes the restricted right-arm rotation distance between S_1 and T_2, came from the word α in $\{x_0^{\pm 1}, x_1^{\pm 1}\}$. We know from Fordham's method of calculating word length with respect to the generating set $\{x_0, x_1\}$ directly from the tree pair diagram that each pair of nodes contributes a certain number of generators to this word length. Fordham calls this the *weight* of the pair of nodes. We see from Fordham's table of weights [9] that any pair of nodes in which one node is a right node has a weight of at most three. So using an extra generator of the form x_0^{-1} to transform this right node into a left node means that these nodes contribute at most four generators each to the length of the word realizing the restricted right-arm rotation distance between T_1 and T_2. We have thus shown the existence of the upper bound

$$d^{\mathcal{T}}_{RRA}(T_1, T_2) \leq 4n - 8.$$

Since rotation distance can only decrease when additional rotations are permitted, this extends immediately to show

$$d^{\mathcal{S}}_{RRA}(T_1, T_2) \leq 4n - 8.$$

Case 2: The tree pair diagram (T_1, T_2) is not reduced.

In this case, since $d^{\mathcal{S}}_{RRA}(T_1, T_2)$ is defined, we know from Lemma 3.6 that there is a partially reduced form of w obtained by applying the usual reduction rules but without reducing instances of x_0 and x_0^{-1} with no $x_1^{\pm 1}$ from the normal form. The number of nodes in the tree pair diagram may reduce to $n' < n$. The proof of this case is now identical to that of Case 1. This produces an upper bound of $4n' - 8 < 4n - 8$ on the restricted right-arm rotation distance between the two trees. □

We now show that the multiplicative constant of 4 is necessary for the above inequality.

Theorem 3.9. *Let $\mathcal{S} = \{x_0, x_{i_1}, \cdots, x_{i_m}\}$. Then there exist trees T_1 and T_2, each with n nodes, so that $d^{\mathcal{S}}_{RRA}(T_1, T_2)$ is defined, and with*

$$4n - 2i_m + 4 \leq d^{\mathcal{S}}_{RRA}(T_1, T_2).$$

The generating set $\mathcal{S}$ used in Theorem 3.9 corresponds to a series of rotations along the right side of the tree from levels 0 to i_m but does not necessarily include all rotations at levels within this range. We now enlarge our generating set to correspond to all rotations at levels 0 to m_i, and work with this set $\mathcal{S}'$ in Theorem 3.10. It will be enough to use this larger set of generators and show that $d^{\mathcal{S}'}_{RRA}(T_1, T_2) \geq 4n - 2i_m + 4$. Thus we prove the following theorem.

Theorem 3.10. *Let $\mathcal{S}' = \{x_0, x_1, x_2, \cdots, x_m\}$. Then there exist trees T_1 and T_2, each with n nodes, so that $d^{\mathcal{S}'}_{RRA}(T_1, T_2)$ is defined, and with*

$$4n - 2m + 4 \leq d^{\mathcal{S}'}_{RRA}(T_1, T_2).$$

The elements we will use to prove this theorem have normal form $x_{m+2}x_{m+3} \cdots x_n x^{-1}_{n-1} \cdots x^{-1}_{m+1}$, and tree pair diagram which we denote (T_1, T_2). These elements are pictured in Figure 7.

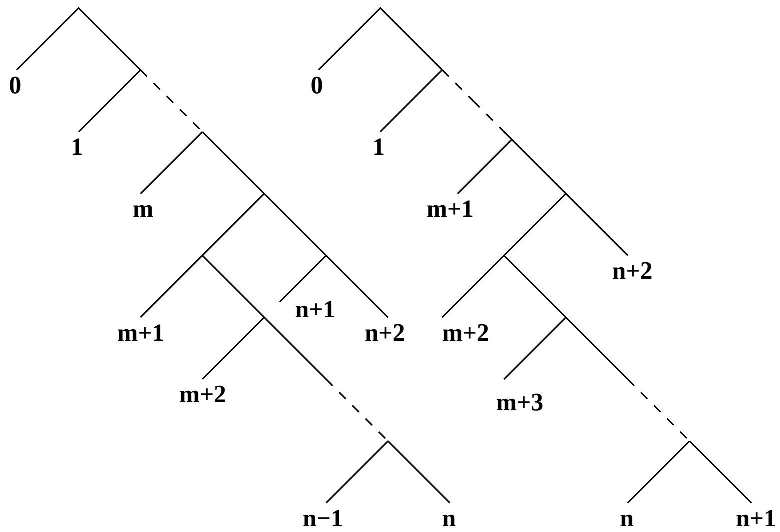

FIGURE 7. The tree pair diagram (T_1, T_2) for words of the form $x_{m+2}x_{m+3}...x_n x^{-1}_{n-1}...x^{-1}_{m+1}$.

Note that Fordham's method for computing exact word length is only valid for the generating set $\{x_0, x_1\}$, so we bound the lengths of these elements indirectly in a series of lemmas by considering the number of rotations needed to create and destroy particular sibling pairs in the tree pair diagram.

We write $[i, i+1]$ if leaves i and $i+1$ form a sibling pair. Performing a rotation corresponding to the generator x_n on a tree T is equivalent to taking the product of x_n with the tree pair diagram $(T, *)$, where $*$ is the tree consisting only of the root node and a series of right nodes. By analyzing the effect of a rotation at a right node on a tree T, we see that there are only three configurations of T which allow a sibling pair to be created or destroyed. These are presented in Figure 8, where capital letters refer to nonempty subtrees of T and lower case letters denote leaf numbers. Right rotation at the appropriate node N along the right side of the tree has the following effect on the sibling pairs.

(i) The pair [a,b] is destroyed and the pair [b,c] is created.
(ii) The pair [a,b] is destroyed.
(iii) The pair [b,c] is created.

We can similarly consider left rotation at the node N, in which case we refer to Figure 9. Left rotation at node N along the right side of the tree has the following effect on the sibling pairs.

(i) The pair [a,b] is created and the pair [b,c] is destroyed.
(ii) The pair [a,b] is created.
(iii) The pair [b,c] is destroyed.

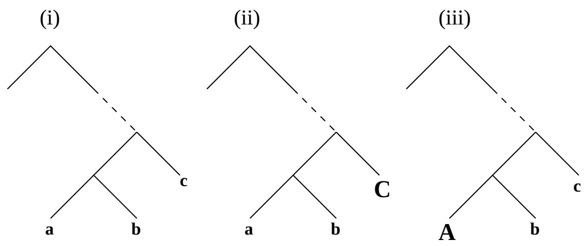

FIGURE 8. Instances where right rotation at node N along the right side of the tree creates or destroys sibling pairs in the tree T. Capital letters represent nonempty subtrees of T, and lower case letters denote leaf numbers.

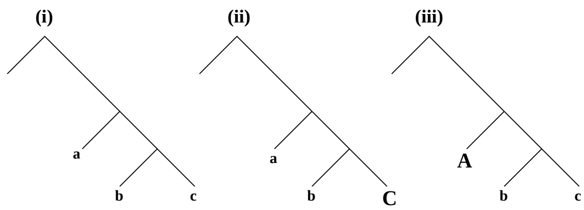

FIGURE 9. Instances where left rotation at node N along the right side of the tree creates or destroys sibling pairs in the tree T. Capital letters represent nonempty subtrees of T, and lower case letters denote leaf numbers.

From these observations, we can see immediately at which nodes it is possible to create and destroy sibling pairs with a set of rotations.

Lemma 3.11. *Suppose a tree T' is obtained from a tree T by applying a right rotation at a node N at level n on the right side of T. If leaves m and $m+1$ are siblings in T and are not siblings in T', then leaves m and $m+1$ are the leaves of an exposed node whose parent is node N. Similarly, if we have the sibling pair $[m, m+1]$ in T' but not in T, then $[m, m+1]$ must be the rightmost node in T' and m must be a leaf in T whose parent is the node at level n in T. Similarly, for left rotations we have the opposite conditions.*

We note that when N is the root node, the only sibling pairs that can be affected by rotation at N consist of the first and last two leaves in the tree.

When we consider the trees T_1 and T_2 in Figure 7 from the word $x_{m+2}x_{m+3}...x_n x_{n-1}^{-1}...x_{m+1}^{-1}$, we see that in T_1, leaves $n-1$ and n are siblings and in T_2, n and $n+1$ are siblings.

Now we consider the number of applications of generators needed to change the sibling pairings from $[n-1, n]$ to $[n, n+1]$ via a minimal length sequence of transformations, expressed as a word $w = g_1 g_2 \dots g_l$. In order to change the pairings, we will need to destroy the sibling pair $[n-1, n]$ and then create the sibling pair $[n, n+1]$. The exposed nodes with siblings $[n-1, n]$ and $[n, n+1]$ are deeply buried in the sense that many rotations are required to affect those nodes and thus those leaf pairings. We measure this depth more precisely with the following definition.

Definition 3.12. *Let c be an exposed node, and define $G(c)$ to be the node connected to c by a maximal length path of downward directed right edges. We also define the ordered pair $D(c) = (p, r)$, where p is the length of the maximal path of right edges from $G(c)$ to c and r is the level in the tree of the ancestor closest to $G(c)$ which is a right caret.*

For example, in the tree T_1 for w and c the node with exposed leaves $[n-1, n]$, we have that $G(c)$ is the node with left leaf $m+1$ and the closest right ancestor of $G(c)$ is its parent. Thus $D(c) = (n-m-2, m+1)$. For the node c' in the tree T_2 with leaves $[n, n+1]$, we have that $G(c')$ is the node with left leaf $m+2$ and $D(c') = (n-m-2, m+2)$.

Lemma 3.13. *Let $w = (T_1, T_2) \in F$ have normal form*

$$x_{m+2}x_{m+3} \cdots x_n x_{n-1}^{-1} \cdots x_{m+1}^{-1}$$

where $n > m+1$. The tree T' resulting from the application of at most $2n-2m-3$ rotations at locations at levels 0 to m along the right side of the tree to T will contain the sibling pair [n-1,n].

Proof. We note that by Lemma 3.11, sibling pairs can be destroyed only when they are connected by a single left edge to a node on the right side of the tree at level m or less, or are the rightmost node in the tree. The exposed node c with leaves $n-1$ and n cannot be moved to be the rightmost node of the tree, as all rotations preserve the natural infix order on the nodes. Thus, until the node c is connected by a single left edge to the right side of the tree at level m or less, leaves $n-1$ and n will remain sibling pairs. We use the ordered pair $D(c)$ to monitor the position of node c while rotations are performed on the tree. The leaves $n-1$ and n will remain sibling pairs until $D(c) = (0, l)$ for some $l \leq m$.

We consider the sequence of trees $S_0 = T_1$, $S_1, S_2, \cdots, S_k$ resulting from performing a series of k rotations corresponding to a sequence of k generators $g_1 g_2 \cdots g_k$. Each S_i is the result of applying g_i to S_{i-1}. We trace the images of the node c through this sequence and denote its image in S_i by c_i. While the exposed leaves of each c_i have the same leaf numbers in S_i, the entries in $D(c_i)$ may change as a result of each rotation.

First we consider the possible changes in $D(c_i) = (r_i, s_i)$ resulting from performing a rotation to the tree corresponding to a single generator. Recall that s_i is the level in the tree of the closest ancestor of $G(c_i)$ on the right side of the tree. There are several cases depending upon the relative location of the node where the rotation takes place and the closest ancestor of $G(c)$ on the right side of the tree.

Here we describe the changes in (r_i, s_i) caused by rotation at a right node at level k, with $0 \leq k \leq m$.

- *Case $k < s_i$:* In this case, the rotation takes place at a level above both c_i and the closest ancestor of $G(c_i)$ on the right side of the tree and we obtain $D(c_{i+1}) = (r_i, s_i - 1) = (r_{i+1}, s_{i+1})$ for a left rotation and $D(c_{i+1}) = (r_i, s_i + 1) = (r_{i+1}, s_{i+1})$ for a right rotation. Note that right rotation may not be possible at level k if the left subtree of the relevant node is empty.
- *Case $k = s_i$:* In this case, a right rotation moves $G(c_i)$ or one of its ancestors which is not on the right side of the tree to the right side of the tree. If $G(c_i)$ is only a single edge away from the right side of the tree and we rotate rightwards, this will change one of the right edges in its path to c_i to a left edge, so $G(c_{i+1})$ is distinct from $G(c)$ and its closest ancestor on the right side of the tree will be one level further from the root node of the tree, so we have $D(c_{i+1}) = (r_i - 1, s_i + 1) = (r_{i+1}, s_{i+1})$. If $G(c_i)$ is more than a single edge away from the right side of the tree, a right rotation will not change $D(c_i)$ and we have $D(c_{i+1}) = D(c_i)$. A left rotation, if permitted, moves $G(c)$ one level further from the root node, but

will not change the location of the closest ancestor of $G(c)$ on the right side of the tree, and we obtain $D(c_{i+1}) = D(c_i)$.

- *Case* $k > s_i$*:* In this case, the right or left rotation takes place further down the right side of the tree and does not affect the location of c_i or $G(c_i) = G(c_{i+1})$, so we have $D(c_{i+1}) = D(c_i)$.

We note that it is possible to move c and $G(c)$ by a sequence of rotations to hang from the left side of the tree, but that they will need to be moved back to the hang from the right side of the tree to split the sibling pair $[n-1, n]$ so this will not happen in any minimal length transformation.

We know that $D(c)$ begins at $(n-m-2, m+1)$, and the sibling pair $[n-1, n]$ is not destroyed until after $D(c_i) = (0, l)$, with $l \leq m$. The only rotations which potentially reduce the first coordinate $n-m-2$ are from the case where $s = k$, with $k \leq m$. There will need to be at least $n-m-2$ such reductions, and each such reduction will increase the second component by 1. We will require at least $n-m-2$ additional reductions to reduce the second coordinate back to its starting value, which do not change the first coordinate. We then must perform at least one additional rotation to decrease the second coordinate to m before the sibling pair in question can be destroyed. This gives a minimum of $2n-2m-3$ rotations before we are in a position to destroy the sibling pair $[n-1, n]$. □

Similarly, in the following lemma we consider the number of rotations required to destroy the sibling pair $[n, n+1]$ in the positive tree T_2 of w.

Lemma 3.14. *Let* T_2 *be the positive tree for the word*

$$x_{m+2}x_{m+3}\cdots x_n x_{n-1}^{-1}\cdots x_{m+1}^{-1}$$

with $m < n$*. The tree* T' *resulting from the application of at most* $2n-2m-2$ *rotations at locations on the right side of the tree at levels less than or equal to* m *will contain the sibling pair [n,n+1].*

Proof. We note that in this case, when c is the node with exposed leaves numbered n and $n+1$, we have $D(c) = (n-m-2, m+2)$ and to reduce $D(c)$ to $(0, l)$ with $l \leq m$ will take at least $(n-m-2)+(n-m-2)+2 = 2n-2m-2$ rotations by the same analysis as in Lemma 3.13. □

These lemmas show that we can find elements with n nodes whose restricted right arm rotation distance satisfies both upper and lower bounds with a multiplicative constant of 4, proving Theorem 3.10.

Proof of Theorem 3.10. We consider the reduced tree pair diagram (T_1, T_2) corresponding to the element $w = x_{m+2}x_{m+3}...x_n x_{n-1}^{-1}...x_{m+1}^{-1}$, as above. Lemma 3.13 shows that any application of $2n-2i_m-3$ allowed rotations to T_1 will still result in a tree with leaves $n-1$ and n paired, and

Lemma 3.14 shows that any application of $2n - 2i_m - 2$ allowed rotations to T_2 will result in a tree with leaves n and $n+1$ paired. Since it will take at least one additional rotation to destroy each sibling pair, we have that the restricted right arm rotation distance between the two trees is at least $4n - 4m - 4$. □

Theorem 3.10 gives a family of pairs of trees with n nodes which satisfy a lower bound on restricted right-arm rotation distance with respect to a generating set $\mathcal{S}'$ which includes all generators from x_0 to x_m. Restricting the generating set to a subset $\mathcal{S}$ of $\mathcal{S}'$ which includes x_0 can only increase the restricted right-arm rotation distance between two trees or cause it to be undefined. In the case of the words used in the proof of Theorem 3.10, the restricted right-arm distance will still be defined by Lemma 3.6, as the trees involved do not involve generators of index lower than the largest level where rotation is allowed. Thus we have proven Theorem 3.9 as well.

4. Bounding right-arm rotation distance

The original arguments of Culik and Wood [8] which give a bound on ordinary rotation distance apply to right-arm rotation distance as well. Their argument is that any binary tree T with n nodes can be transformed to or from the all-right tree with n nodes by no more than $n-1$ rotations, all of which can be chosen to lie on the right arm of the tree. Thus, the right-arm rotation distance between two trees T_1 and T_2 each with n nodes is no more than $2n-2$, as we can transform T_1 to the all-right tree and from there transform it to T_2. While this bound is not optimal for the original rotation distance, we show that it is optimal for right-arm rotation distance.

Theorem 4.1. *For each $n \geq 3$, there are rooted binary trees T_1 and T_2 each with n nodes so that the right-arm rotation distance between them satisfies $d_{RA}(T_1, T_2) = 2n - 2$.*

Proof. To prove this we consider the elements of F with normal form $x_0x_1x_2x_3 \dots x_{n-2}x_{n-3}^{-1}x_{n-4}^{-1} \dots x_1^{-1}x_0^{-2}$, pictured in Figure 10, which have n nodes and have word length $2n-2$ with respect to the infinite generating set for F. It follows from Proposition 3.7 that this is also the right-arm rotation distance between the two trees. □

5. Left-arm and spinal rotation distances

We now consider rotation distances which include rotations at nodes along the left side of the tree instead of or in addition to nodes along the right side of the tree. It is clear by symmetry that *restricted left-arm rotation distance*, which allows rotations at only nodes on the left side of the tree and

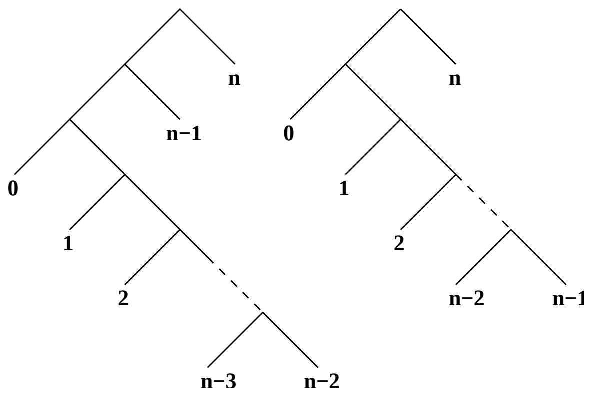

FIGURE 10. The tree pair diagram for words of the form $x_0x_1x_2x_3\ldots x_{n-2}x_{n-3}^{-1}x_{n-4}^{-1}\cdots x_1^{-1}x_0^{-2}$, with n nodes and length $2n-2$ with respect to the infinite generating set.

the root node, will satisfy the same bounds as restricted right-arm rotation distance. Similarly, *left-arm rotation distance*, which allows rotations at any node along the left arm of the tree, will satisfy the sharp upper bound of $2n-2$ on trees with n nodes.

Finally, we consider a rotation distance which allows rotations at the root node, a finite nonempty collection of nodes on the right side of the tree, and at a finite nonempty collection of nodes on the left side of the tree. Since all nodes where rotations are permitted lie on the spine of the tree, we call such a rotation distance a *restricted spinal rotation distance.* Again, though allowing rotations at finitely many locations on both the right and left arms of the tree may reduce the rotation distance between some pairs of trees, the multiplicative constant of 4 in the bound does not decrease. In terms of Thompson's group F, rotation at level n on the left arm of the tree can be expressed as $y_n = x_0^n x_1 x_0^{-n-1}$.

Theorem 5.1. *Let $\mathcal{S} = \{x_0, x_{i_1}, \cdots, x_{i_m}, y_{j_1}, \cdots, y_{j_l}\}$ where x_i is a generator of F and $y_n = x_0^n x_1 x_0^{-n-1}$. Then there exist trees T_1 and T_2 with n nodes for which $d_{RRA}^{\mathcal{S}}$ is defined that satisfy*

$$4n - 2i_m + 4 \leq d_{RRA}^{\mathcal{S}}(T_1, T_2).$$

Proof. We use the trees T_1 and T_2 from Figure 7 again to establish this bound. We note that analogues of Lemmas 3.13 and 3.14 hold, as the additional allowed rotations on the left side of the tree have the same effect on the sibling pairs as the right rotations and thus on $D(c)$, so the argument is analogous. $\square$

References

[1] Kenneth S. Brown. The geometry of rewriting systems: a proof of the Anick-Groves-Squier theorem. In *Algorithms and classification in combinatorial group theory (Berkeley, CA, 1989)*, pages 137–163. Springer, New York, 1992.

[2] Kenneth S. Brown and Ross Geoghegan. An infinite-dimensional torsion-free FP_∞ group. *Inventiones mathematicae*, 77:367–381, 1984.

[3] J. W. Cannon, W. J. Floyd, and W. R. Parry. Introductory notes on Richard Thompson's groups. *Enseign. Math.*, 42(3-4):215–256, 1996.

[4] Sean Cleary. Restricted rotation distance between binary trees. *Inform. Process. Lett.*, 84(3):333–338, 2002.

[5] Sean Cleary and Jennifer Taback. Restricted rotation distance. *Information Processing Letters*, 88(5):251–256, 2003.

[6] Sean Cleary and Jennifer Taback. Thompson's group F is not almost convex. *J. Algebra*, 270(1):133–149, 2003.

[7] Sean Cleary and Jennifer Taback. Combinatorial properties of Thompson's group F. *Trans. Amer. Math. Soc.*, 356(7):2825–2849 (electronic), 2004.

[8] Karel Culik II and Derick Wood. A note on some tree similarity measures. *Inform. Process. Lett.*, 15(1):39–42, 1982.

[9] S. Blake Fordham. Minimal length elements of Thompson's group F. *Geom. Dedicata*, 99:179–220, 2003.

[10] Donald E. Knuth. *The Art of Computer Programming. Volume 3.* Addison-Wesley, Reading, Mass, 1973. Sorting and searching.

[11] Jean Pallo. An efficient upper bound of the rotation distance of binary trees. *Inform. Process. Lett.*, 73(3-4):87–92, 2000.

[12] Jean Marcel Pallo. Right-arm rotation distance between binary trees. *Inform. Process. Lett.*, 87(4):173–177, 2003.

[13] R. Rogers. On finding shortest paths in the rotation graph of binary trees. In *Proc. Southeastern Int'l Conf. on Combinatorics, Graph Theory, and Computing*, volume 137, pages 77–95, 1999.

[14] Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1(3):647–681, 1988.

Department of Mathematics, The City College of New York & The CUNY Graduate Center, New York, NY 10031

E-mail address: cleary@sci.ccny.cuny.edu

Department of Mathematics, Bowdoin College, Brunswick, ME 04011

E-mail address: jtaback@bowdoin.edu