Bounding restricted rotation distance

Sean Cleary [1]

Department of Mathematics, City College of CUNY, New York NY 10031

Jennifer Taback [2]

Department of Mathematics and Statistics, University at Albany, Albany, NY 12222

Abstract

Restricted rotation distance between pairs of rooted binary trees quantifies differences in tree shape. Cleary exhibited a linear upper bound of $12n$ for the restricted rotation distance between two trees with n interior nodes, and a lower bound of $(n-1)/3$ if the two trees satisfy a reduction condition. We obtain a significantly improved sharp upper bound of $4n-8$ for restricted rotation distance between two rooted binary trees with n interior nodes, and a significantly improved sharp lower bound of $n-2$, again with the requirement that the trees satisfy a reduction condition. These improvements use work of Fordham to compute the word metric in Thompson's group F.

Key words: algorithms and data structures, binary trees, rotation distance

1 Introduction

In this paper, we consider several methods of quantifying the difference in "shape" between two rooted binary trees of the same size. In each method, this is done by counting the number of elementary changes, or *rotations* necessary to transform one tree to the other. Depending on where in the tree these rotations are allowed to occur, one obtains either the rotation distance, or restricted rotation distance between the two trees.

When rotations are allowed at any node of the tree, we obtain *rotation distance*, analyzed by Sleator, Tarjan and Thurston in [1]. Given two trees with n interior nodes each, they obtain a bound of $2n-6$ on the number of rotations needed to transform one tree to the other. Moreover, they show that this bound is the best possible. Polynomial time algorithms of Pallo [2] and Rogers [3] estimate rotation distance, but no efficient algorithm is known to compute it exactly.

When rotations are allowed only at the root node and the right child of the root, we obtain *restricted rotation distance* d_{RR}. This restriction allows connections with Thompson's group F, discussed below, which is generated by such rotations and allows efficient

[1] supported by PSC-CUNY grant #63438-0032

[2] partial support from an NSF-AWM Mentoring Travel grant.

estimation of restricted rotation distance, as described in [4].

The trees considered below are all ordered, rooted binary trees with n interior nodes, where each interior node has 2 children. These trees are sometimes called *extended binary trees* [5] or *0-2 trees*, and we will refer to them as *binary trees*. Additionally, we use *node* to mean an interior node, and *leaf* for a non-interior node. We number the nodes of our tree using the infix ordering.

Cleary explores the relationship between bounding restricted rotation distance and Thompson's group F in [4]. This group is an abstract group which has been studied extensively in the fields of combinatorial group theory, measure theory and logic; an overview of F is given by Cannon, Floyd and Parry in [6]. Using metric estimates in this group, Cleary obtains an upper bound of $12n$ on the restricted rotation distance between two binary trees with n nodes. If the trees satisfy a reduction condition, he obtained a lower bound of $(n-1)/3$ on the restricted rotation distance.

In this paper, we apply a process of Fordham [7] which computes word length in F to improve Cleary's bounds on restricted rotation distance. In particular, we obtain a sharp upper bound of $4n-8$ on the the restricted rotation distance between two trees with n interior nodes. We note that restricting the number of permitted rotations from n to two only increases the $2n-6$ upper bound on ordinary rotation distance by a factor which asymptotically approaches 2.

2 Background

We now make precise the notion of rotation at a node in a tree, pictured as the two trees in Fig. 1. *Right rotation* at the node N transforms the left tree in Fig. 1 into the right tree

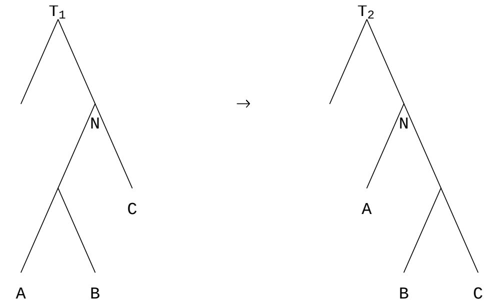

Fig. 1. Right rotation at node N. The letters A, B and C represent arbitrary subtrees of T_1 and T_2. Left rotation at node N is defined analogously.

in the figure. *Left rotation* is defined to be the inverse operation, transforming the right tree to the left one. We always use T_1 and T_2 to denote trees with the same number of leaves.

When computing ordinary rotation distance, we allow the node N to be any node of the tree; for restricted rotation distance, we restrict N to either the root node or the right child of the root.

Restricted rotation distance has the potential to be much greater than ordinary rotation distance, since many rotations at these two distinguished nodes may be needed to accomplish the equivalent of a single rotation at a node which is far from the root node. Thompson's group F is useful in understanding and computing restricted rotation distance.

3 Thompson's Group F and Tree Pair Diagrams

Thompson's group F can be defined using combinatorial, analytic or geometric methods, described by Cannon, Floyd and Parry in [6].

The group F has a presentation with an in-

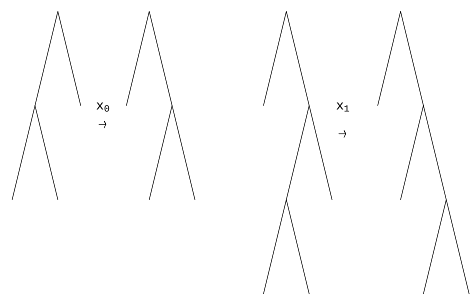

Fig. 2. The generators x_0 and x_1 of F, as tree pair diagrams.

finite number of generators and relations:

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

Group elements have standard normal forms, described by Brown and Geoghegan in [8] and the relations provide an efficient method for rewriting words into normal form.

Thompson's group F also has a finite presentation given only in terms of the generators x_0 and x_1. Powers of x_0 conjugate x_1 to x_n; thus we can generate the entire group using those two elements. All of the infinitely many relations are consequences of two relations of length 10 and 14 (see [8]).

The group F is defined geometrically using *tree pair diagrams*, each of which is a pair of rooted binary trees with the same number of leaves. Examples and a description of the equivalence of these two interpretations is given in [9]. The generators x_0 and x_1 are represented by the tree pair diagrams given in Fig. 2. We can view x_0 as performing a right rotation at the root caret, and x_1 as performing a right rotation at the right child of the root.

We say that a tree pair diagram (T_1, T_2) is *unreduced* if both T_1 and T_2 contain a node with two leaves numbered i and $i+1$. Any tree pair diagram which is not unreduced is called *reduced*. Every element of F corresponds to a unique reduced tree pair diagram. An example of transforming an unreduced tree pair diagram into a reduced one is given in Fig. 3 of [4]. We view F as a group of tree pair diagrams with the operation of composition. It is sometimes necessary to expand the trees, creating unreduced representatives of elements of F, in order to perform this composition (see [6]).

4 The Metric on F and Restricted Rotation Distance

Let G be a group and $X = \{x_1, x_2, \cdots x_n\}$ a finite set of generators for G. There is a canonical way of associating a geometric picture to the pair (G, X), called a *Cayley graph*. Namely, the Cayley graph $\Gamma(G, X)$ has as vertices all elements of G, with a directed edge from g to $g \cdot x_i$, for $g \in G$ and $x_i \in X$. This graph is a metric space using the *word metric*, which defines the distance between elements g and h as the minimal length of a path between them in $\Gamma(G, X)$. For $g \in G$, the length of g is the distance from g to the identity in this word metric. An introduction to Cayley graphs and other objects in geometric group theory is given in Epstein *et al* [10].

Given $f = (T_1, T_2) \in F$, a representation of f as a product of the generators $x_0^{\pm 1}$ and $x_1^{\pm 1}$ with m terms can be thought of as a sequence of m rotations at the root and right child of the root which transforms the tree T_1 into the tree T_2. Thus, the restricted rotation distance between T_1 and T_2 will be no more than m. There may be alternate ways to express f as a product of $x_0^{\pm 1}$ and $x_1^{\pm 1}$ with fewer total terms, so unless we know that the expression of f as a product of m generators is minimal, we obtain only an upper bound for the restricted rotation distance.

In [4], Cleary estimated restricted rotation distance by estimating the word metric in F. Fordham [7] developed a method for com-

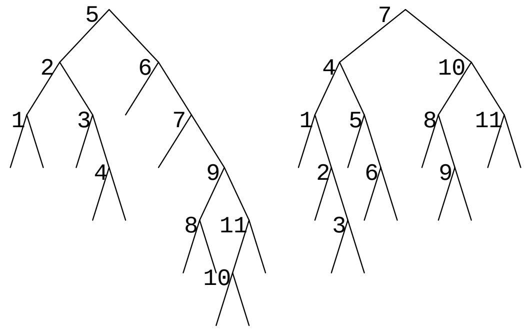

Fig. 3. The reduced tree pair diagram for the word $w = x_0^2 x_1 x_2 x_4 x_5 x_7 x_8 x_9^{-1} x_7^{-1} x_3^{-1} x_2^{-1} x_0^{-2}$ with nodes numbered according to the infix method. The correspondence between the word and the trees is obtained using the notion of *leaf exponents*, see [7] or [9].

puting the exact length of an element of F, which we use here to compute the exact restricted rotation distance between two binary trees. Given T_1 and T_2, the length of the element $f = (T_1, T_2)$ in F is exactly the restricted rotation distance between T_1 and T_2, and a minimal length representative of f gives a sequence of rotations which achieves that distance.

5 Fordham's calculation of word length in F

Fordham [7] presents an algorithm for finding the exact length of an element of F given by a tree pair diagram. His method orders the interior nodes of the trees, classifies them into seven types, and assigns weights listed in a table.

First, we note that the interior nodes of a tree T can be given the natural infix order where the nodes in the left subtree of a node precede it and the nodes in the right subtree follow it. Thus, we number nodes from 1, the leftmost node, to n, the rightmost one, as used in the labelling of the nodes in the trees in Fig. 3.

Once ordered, nodes are classified first into several basic categories. *Left nodes* are the root node and the nodes connected to the root by paths consisting entirely of left edges. *Right nodes* are the nodes connected to the root by paths consisting entirely of one or more right edges. The left nodes in Fig. 3 are 1, 2 and 5 on the left tree and 1, 4 and 7 on the right tree. The right nodes are 6, 7, 9 and 11 and 10 and 11, respectively.

Nodes are further classified as being exactly one of the following *node types*:

- L_0: the node numbered 1, necessarily a left node.
- L_L: left nodes not numbered 1
- R_I: right nodes whose immediate successor node, in the infix ordering, is not a right node.
- R_0: right nodes whose successors in the infix ordering are all right nodes.
- R_{NI}: right nodes not of type R_I or R_0.
- I_0: nodes which are neither left nor right and have no right child.
- I_R: nodes which are neither left nor right and have a right child.

The node type of the node numbered i in a tree T is denoted $\tau_i(T)$. We define a *node pairing* for each node pair in the tree pair diagram (T_1, T_2) as the pairing of types $(\tau_i(T_1), \tau_i(T_2))$. For example, as node 1 of a nonempty T_1 will always be a node of type L_0, as will the node 1 in T_2, we have $(\tau_1(T_1), \tau_1(T_2)) = (L_0, L_0)$.

As an example, for the tree pair diagram in Fig. 3, we see that the nodes in the left tree T_1 are, in infix order, of types $L_0, L_L, I_R, I_0, L_L, R_{NI}, R_I, I_0, R_I, I_0$ and R_0. The nodes in the right hand tree T_2 of Fig. 3, in infix order, are of types $L_0, I_R, I_0, L_L, I_R, I_0, L_L, I_R, I_0, R_0$ and R_0.

Fordham [7] defines the *weight* of a node pairing as follows: for the node pairing (L_0, L_0), the weight is zero; all other node pairing weights are prescribed according to the following table of weights.

weight	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 proves the following theorem to calculate word lengths in F, with respect to the word metric arising from the finite generating set $\{x_0, x_1\}$ of F. In the arguments below, we are always interested in word length with respect to this finite generating set.

Theorem 1 (Theorem 2.5.1 [7]) *Given an element $f \in F$ with reduced tree pair diagram (T_1, T_2), the length of f with respect to the generating set $\{x_0^{\pm 1}, x_1^{\pm 1}\}$ is given by the sum of the weights of the node pairings of (T_1, T_2).*

Considering the word w in Fig. 3, we see that the nodes numbered one in each tree have type pairing (L_0, L_0), which has weight 0. The nodes numbered 2 have types (L_L, I_R) which contributes 2 to the length 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 is 20. That means that there is an expression for w of length 20 in terms of $x_0^{\pm}, x_1^{\pm}$ and there are no expressions for w of length less than 20. That sequence of 20 generators, when thought of as rotations applied to the left-hand tree of Fig. 3 will transform it to the right-hand tree in that figure. There may be other sequences of length 20 or more which accomplish the same transformation (for this particular example, there are exactly two sequences of length 20), but there are none of shorter length.

We use this process for finding word lengths to get an upper bound of $4n - 8$ for the rotation distance between two trees with n nodes, which is sharper than the $4n - 4$ obtained in Thm. 3.1 in Cleary and Taback [9].

Theorem 2 *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 $d_{RR}(T_1, T_2) \leq 4n - 8$.*

Proof: Consider the element f of F given by the tree pair diagram (T_1, T_2). Such a tree pair diagram may be unreduced. If this is the case, we form the reduced tree pair diagram which will involve fewer nodes and leaves. We know that (T_1, T_2) has a total of at most n node pairs, one of which is the first node pair of type (L_0, L_0) and thus of weight 0. The remaining $n-1$ node pairs have weights at most 4, so the total weight of the tree pair diagram is at most $4n - 4$, as obtained in [9]. Thus, f can be represented by a string of generators $x_0^{\pm 1}$ and $x_1^{\pm 1}$ of length no more than $4n - 4$. Since these generators correspond to left and right rotations at the root and right child of the root, it is possible to transform T_1 to T_2 with the sequence of rotations corresponding to this string of generators. Thus, we have a simple upper bound for rotation distance. To improve this bound to $4n - 8$, we consider the following. Each pair of nodes in the tree pair diagram contributes at most 4 to the word length, since that is the largest value in the table of node pair weights. We focus on computing the deficit of word length from $4n$.

First, we notice every tree pair has a pair of nodes of types (L_0, L_0) which has weight 0. As mentioned above, this pair of nodes contributes a deficit of 4 to the total weight deficit. The final node in a tree is either the root node and thus of type L_L, or a node on the right side of the tree, necessarily of type R_0. We consider the three possible cases for the final node pair.

Case A: The final node in both T_1 and T_2 is of type L_L. Since these nodes must be paired, there is a pairing of type (L_L, L_L) which has weight 2 and thus the deficit for that pair of nodes is 2, giving a total deficit of at least 6. We now consider the penultimate node pairing. If the last node is of type L_L, it would necessarily be the root and the next-to-last node could only be of type L_L or I_0. Note that both penultimate nodes cannot be of type I_0 as that would yield an unreduced representative of f. Each of these possible penultimate node pairings has weight at most 2, giving a total deficit of at least 8.

Case B: The final nodes in T_1 and T_2 have types L_L and R_0, although not necessarily respectively. Since these nodes must be paired, we have a pairing of type (L_L, R_0) or (R_0, L_L) which has weight 1 and thus the deficit for that node pair is 3, giving a total deficit of at least 7. If the final caret of a tree is of type L_L, the penultimate caret may be of type L_L or I_0. If the final caret of a tree is of type R_0, the penultimate caret may be of type I_0, R_0, or L_L. All possible caret pairings for the penultimate caret have weight 2 or less, yielding a deficit of at least 9.

Case C: Both trees have final nodes of type R_0. Since these nodes must be paired, we have a pairing of type (R_0, R_0) which has weight zero and thus the deficit for that pair of nodes is 4, giving a total deficit of at least 8. Note that, similar to case B, there must be an additional deficit from the penultimate node pairing, so the total deficit will in fact be 10 or greater.

Thus the deficit is at least 8, giving an upper bound on word length, and hence on restricted rotation distance, of $4n - 8$. □

Theorem 3 *For any $n \geq 3$, there are rooted binary trees T_1 and T_2 each with n nodes such that the restricted rotation distance $d_{RR}(T_1, T_2)$ between them is $4n - 8$.*

Proof: We can realize this bound in several possible ways. Perhaps the simplest word which realizes this bound for $n \geq 4$ is $x_0x_1x_2x_3 \ldots x_{n-2}x_{n-3}^{-1}x_{n-4}^{-1} \ldots x_1^{-1}x_0^{-2}$ pictured in Fig. 4. This tree pair diagram has n nodes in each tree and represents a word of length $4n - 8$.

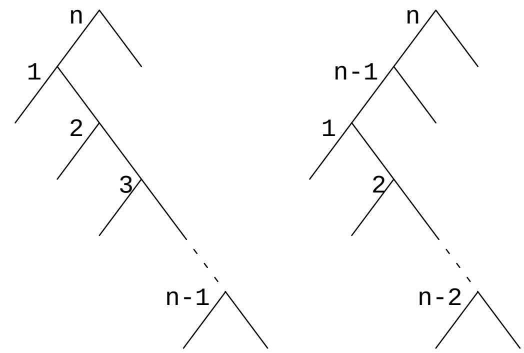

Fig. 4. The tree pair diagram for words of the form $x_0x_1x_2x_3 \ldots x_{n-2}x_{n-3}^{-1}x_{n-4}^{-1} \ldots x_1^{-1}x_0^{-2}$, with the nodes numbered according to the infix method.

The pairs of nodes in these trees are of the following types: $(L_0, L_0), (I_R, I_R), \ldots, (I_R, I_R), (I_R, I_0), (I_0, L_L), (L_L, L_L)$. The pairings give deficits $4, 0, 0, \ldots 0, 0, 2, 2$, respectively, with total deficit 8, proving that the bound is sharp. □

We can also use the metric on F to obtain a lower bound on the number of rotations at the root and right child of the root needed to transform T_1 to T_2 improving the estimate in [4] by using the result from [9].

Theorem 4 (Theorem 3.1 [9]) *If T_1 and T_2 are rooted binary trees with n nodes each, which form a reduced tree pair diagram (T_1, T_2), then the restricted rotation distance $d_{RR}(T_1, T_2)$ is at least $n - 2$.*

Proof: We consider pairings which give weight 0. There will be exactly one pairing of type (L_0, L_0) and at most one pairing of type (R_0, R_0), since if there are two such pairings, the tree pair diagram would be unreduced. All of the at least $n - 2$ remaining nodes will contribute at least 1 from the non-zero entries according to the table above. Thus, the weight of the word represented by (T_1, T_2) is at least $n-2$ and so $d_{RR}(T_1, T_2) \geq n-2$ as desired. □

We note that the bound in Theorem 4 is again sharp simply by considering the tree pairs for $x_1^{\pm(n-2)}$ which have n interior nodes and length $n-2$.

References

[1] D. D. Sleator, R. E. Tarjan, W. P. Thurston, Rotation distance, triangulations, and hyperbolic geometry, J. Amer. Math. Soc. 1 (3) (1988) 647–681.

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

[3] 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, Vol. 137, 1999, pp. 77–95.

[4] S. Cleary, Restricted rotation distance between binary trees, Inform. Process. Lett. 84 (6) (2002) 333–338.

[5] D. E. Knuth, The Art of Computer Programming. Volume 3, Addison-Wesley, Reading, Mass, 1973, sorting and searching.

[6] J. W. Cannon, W. J. Floyd, W. R. Parry, Introductory notes on Richard Thompson's groups, L'Ens. Math. 42 (1996) 215–256.

[7] B. Fordham, Minimal length elements of Thompson's group F, Ph.D. thesis, Brigham Young Univ (1995).

[8] K. S. Brown, R. Geoghegan, An infinite-dimensional torsion-free FP_∞ group, Inventiones mathematicae 77 (1984) 367–381.

[9] S. Cleary, J. Taback, Combinatorial properties of Thompson's group, Transactions of the AMS .

[10] D. B. Epstein, J. W. Cannon, D. F. Holt, S. Levy, M. S. Paterson, W. P. Thurston, Word Processing in Groups, Jones and Bartlett, 1992.