Restricted Rotation Distance between Binary Trees

Sean Cleary [1]

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

Abstract

Restricted rotation distance between pairs of rooted binary trees measures differences in tree shape and is related to rotation distance. In restricted rotation distance, the rotations used to transform the trees are allowed to be only of two types. Restricted rotation distance is larger than rotation distance, since there are only two permissible locations to rotate, but is much easier to compute and estimate. We obtain linear upper and lower bounds for restricted rotation distance in terms of the number of interior nodes in the trees. Further, we describe a linear-time algorithm for estimating the restricted rotation distance between two trees and for finding a sequence of rotations which realizes that estimate. The methods use the metric properties of the abstract group known as Thompson's group F.

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

1 Introduction

There are a number of ways of measuring the difference in shape between two rooted binary trees with the same number of leaves. In the particular case of rooted binary trees with a left-to-right ordering on the leaves (such as binary search trees), the difference of shape is of particular interest in balancing of trees. Rotations are a small change in the shape of a binary tree commonly used as primitive steps in tree balancing (see Knuth, [1]). Rotation distance measures this difference by counting the minimal number of fundamental rotations, which can take place at any node, to transform one tree into the other. Work by Culik and Wood [2], Pallo [3,4] and Makinen [5] has clarified the properties of rotation distance. Sleator, Tarjan and Thurston [6], by using geometric methods, obtained an upper bound of $2n-6$ rotations needed to transform one rooted binary tree with n interior nodes into

---

Email address: `cleary@sci.ccny.cuny.edu` (Sean Cleary).

URL: `http://www.sci.ccny.cuny.edu/∼cleary` (Sean Cleary).

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

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. Later work by Luccio and Pagli [7], Pallo [4], Hanke, Ottmann, and Schuierer [8] and Rogers [9] has shown bounds using methods which do not rely on hyperbolic geometry. There is no known polynomial-time algorithm for computing rotation distance, but work of Pallo [4] and Rogers [9] has given polynomial-time algorithms which estimate rotation distance.

We consider ordered, rooted binary trees with n interior nodes and where each interior node has 2 children. Such trees are commonly called *extended binary trees* [1] or *0-2 trees.* In the following, 'binary tree' refers to such a tree, 'node' refers to an interior node, and 'leaf' refers to a non-interior node. Our trees will have $n+1$ leaves numbered in left-to-right order from 0 to n. With the ordinary rotation distance, a rotation can take place at any of the n nodes in the rooted binary tree with $n+1$ leaves. We define restricted rotation distance, where we consider the same rotation operation, but we only allow rotation at one of two nodes– either the root or the right child of the root. (Note this restricted rotation distance is different than in [10], where the rotations are restricted to allow rotations only at nodes with an external edge.) With this minimal sufficient set of allowed rotations, we find an upper bound of $12n$ for the restricted rotation distance between any two rooted trees with n nodes, and a lower bound of $(n-1)/3$ for the minimum number of such rotations needed to change one tree to another if the trees satisfy a reduction condition. We describe a linear-time algorithm to find the rotations to transform the tree. These estimates and algorithm are obtained by considering the word metric on Thompson's group F, an abstract group with many remarkable properties, which has been extensively studied in the fields of combinatorial group theory, measure theory and logic. Cannon, Floyd and Parry [11] have written an excellent introduction to F . The metric properties of F have been studied by Burillo [12] and by Burillo, Cleary and Stein [13], leading to this understanding of restricted rotation distance in the context of rooted binary trees.

2 Restricted Rotation Distance

Rotation at a node of a rooted binary tree is defined as as a simple change to T and is illustrated in Figure 1 as taking the left-hand tree to the right-hand one. Left rotation at a node is the inverse operation, and takes the right hand tree to the left-hand one. The *rotation distance* $d_R(T_1, T_2)$ between two rooted binary trees T_1 and T_2 with the same number of leaves is the minimum number of rotations needed to transform T_1 to T_2, where the rotations can be performed at any node. In the following, T_1 and T_2 are trees with the same number of leaves. Allowing rotations at any node gives a sufficient set of fundamental moves– any tree T_1 can be transformed to any other tree T_2 by a sequence of such rotations. We can consider smaller sets of allowed rotations, occurring at only a subset of all nodes. Culik and Wood [2] showed that allowing rotations only to take place at nodes on the right "spine" is sufficient to transform any T_1 to T_2, and takes no more than $2n - 2$ steps. Their allowed set of rotations grows with the size of the trees. We can restrict further to finite sets of nodes where rotation is allowed. Allowing rotation at only a single node is insufficient to transform any T_1 to T_2 but we can get a minimal sufficient set of fundamental rotations if we choose the allowed two nodes to be the root and the right child of the root. It is not immediately clear that it is possible to transform any T_1 to T_2 by a sequence of rotations taking place at only those two nodes, but that will be shown as Lemma 1 below. With respect to this minimal set of fundamental moves, we define the *restricted rotation distance* $d_{RR}(T_1, T_2)$ between two rooted binary trees T_1 and T_2 with the same number of leaves as the minimum number of rotations at the root or the right child of the root needed to transform T_1 to T_2. We can think of this restricted setting of rotation distance as analogous to the restriction of array operations to stack operations. With ordinary rotation distance, we are able to randomly access any node and perform a rotation there, but with restricted rotation distance, we are limited to access only the root and right child of the root. With this restriction, transformations take more steps, just as using stack operations instead of array operations to accomplish tasks takes more steps. Surprisingly, the number of steps required to transform T_1 to T_2 is still linear in the size of the tree, despite the fact that it may take many rotations even to move a particular node to one of the two places where rotations are allowed.

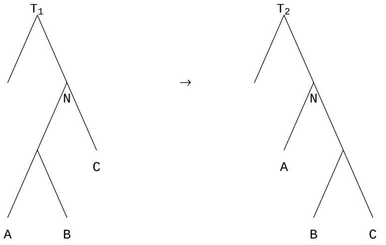

Fig. 1. Right rotation at node N

Clearly, the restricted rotation distance between two trees is bounded below by (ordinary) rotation distance. In general, restricted rotation distance is much greater as it may take many rotations at the two distinguished nodes to accomplish the equivalent operation of a single rotation at a single node at a great distance from the root. There is a trade–off; though the number of steps required to perform the transformation is higher, the individual steps are very simple and only of two types. To understand restricted rotation distance, we study Thompson's group F.

3 Thompson's Group F and Tree Pair Diagrams

Thompson's group F can be defined in three equivalent ways, all of which are useful in understanding the structure of the group. Analytically, we can define F as a group of piecewise-linear homeomorphisms of the unit interval. Combinatorially, we can define F in terms of generators and relations. Geometrically, we can regard F as equivalence classes of pairs of rooted binary trees. The remarkable equivalence of these three characterizations of F is described in [11].

Analytically, we define F as the group of orientation-preserving piecewise-linear homeomorphisms from $[0, 1]$ to $[0, 1]$ 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 infinite presentation:
$\Big\langle x_0, x_1, \ldots | x_i^{-1} x_n x_i = x_{n+1}, \forall i < n \Big\rangle$.
There is a set of normal forms for elements of F given by
$x_{i_1}^{r_1} x_{i_2}^{r_2} \ldots x_{i_k}^{r_k} x_{j_l}^{-s_l} \ldots x_{j_2}^{-s_2} x_{j_1}^{-s_1}$
with $r_i, s_i > 0$, $i_1 < i_2 \ldots < i_k$ and $j_1 < j_2 \ldots < j_l$. This normal form is unique if we further require a reduction condition that when both x_i and x_i^{-1} occur, so does x_{i+1} or x_{i+1}^{-1} as discussed in [14]. The relations provide a quick and efficient manner to rewrite a word into normal form, and form a complete rewriting system, as described in [15].

There is also a finite presentation for F; since x_0 conjugates x_1 to x_2 and similarly higher-index generators are also conjugates of x_1 by x_0, the two generators x_0 and x_1 suffice to generate the whole group. Furthermore, all of the infinitely many relations are consequences of two relations of length 10 and 18 (see [14]).

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 [11]. Essentially, a rooted binary tree can be regarded as a procedure for constructing a subdivision of the unit interval by successive halving of subintervals. A pair of such trees (S, T) gives an element of F from the analytical perspective by considering the element f which is the piecewise-linear homeomorphism which realizes the interpolation of subdivisions described by the source tree S and the target tree T. The equivalence between the rooted tree pair diagram perspective and the combinatorial perspectives described above by the infinite and finite presentations is described in [11], and the tree pair diagrams associated to generators x_0 and x_1 are pictured in Figure 2. Note that the generator x_0 performs a right rotation at the root and that the generator x_1 performs a right rotation at the right child of the root. The inverses x_0^{-1} and x_1^{-1} are left rotations at their respective nodes. The (optional) generators in the infinite generating set defined as $x_n = x_0^{-n+1} x_1 x_0^{n-1}$ are right rotations at the node n levels down on the right side of the tree. The group operation in terms of the tree pair diagram representation is composition; it may be necessary to expand trees (by changing a leaves to an interior nodes and adding new leaves) to perform composition (see [11].) Given any pair of trees (T_1, T_2) with the same number of leaves, we can consider that as representing an element of F. And given any element f of F, we can find (T_1, T_2) representing f, which will be unique if we further require that the tree pair (T_1, T_2) is reduced, as described below. Thus, there is a one-to-one correspondence between reduced tree pair diagrams and elements of F. Now we are in a position to prove the lemma:

Lemma 1 *A rooted binary tree* T_1 *can be transformed into any other tree* T_2 *with the same number of leaves by a sequence of rotations at the root and right child of the root.*

Proof: The tree pair diagram (T_1, T_2) represents an element of F. Since x_0 and x_1 generate F, we can express the rotations needed to change T_1 to T_2 as a sequence of right and left rotations at the root $(x_0^{\pm 1})$ and right child of the root $(x_1^{\pm 1})$. Thus we have the lemma.

We can also prove the lemma by noting that we can express a rotation at any node on the right side of the tree n levels down as the conjugate $x_n = x_0^{-n+1} x_1 x_0^{n-1}$. Since Culik and Wood [2] showed that rotations at nodes on the right side of the tree suffice to change any T_1 to T_2, thus x_0 and x_1 suffice.

4 The Metric on F and Restricted Rotation Distance

Given a group G in terms of a finite presentation, we define the *length* of a word g in G as the number of generators counted with unsigned multiplicity in a minimal-length representative of g as a word in terms of the generators. This length is precisely the distance from the identity in the Cayley graph for the group G with respect to that generating set, where each edge in the Cayley graph is declared to have length 1. The field of geometric group theory has concerned itself with understanding the metric properties of groups and the consequences of metric hypotheses on groups; see Gromov [16] and Epstein *et al.* [17] for an introduction.

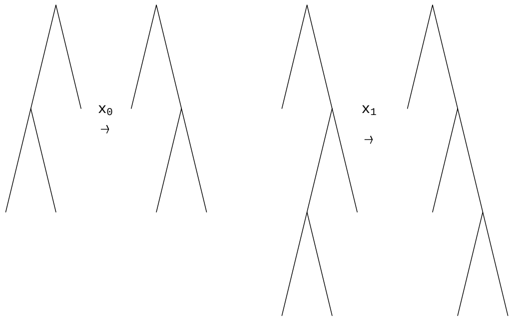

Fig. 2. Elements x_0 and x_1 of F

Burillo [12] and Burillo, Cleary and Stein [13] analyzed the word metric on F. Given a word x in F, there are many representatives of x in terms of the finite generating set $\{x_0, x_1\}$. The length $|x|$ with respect to that finite generating set is the shortest such representative. The relevant result from [13] is

Theorem 1 *Let* $x \in F$ *have normal form* $x_{i_1}^{r_1} x_{i_2}^{r_2} \ldots x_{i_n}^{r_n} x_{j_m}^{-s_m} \ldots x_{j_2}^{-s_2} x_{j_1}^{-s_1}$, *and let* $D(x) = r_1 + r_2 + \ldots + r_n + s_1 + s_2 + \ldots + s_m + i_n + j_m$. *Then we have* $\frac{D(x)}{3} \leq |x| \leq 3D(x)$.

We can define $N(x)$ to be the number of interior nodes in either tree of the reduced tree pair diagram representing x. We showed in [13] that:

Theorem 2 *Let* $x \in F$ *be as above, then* $\frac{D(x)}{4} \leq N(x) \leq D(x)+1$. *Thus we have* $\frac{N(x)-1}{3} \leq |x| \leq 12N(x)$.

These results give us an expression for the maximum rotation distance in terms of the number of nodes:

Theorem 3 *Given two rooted binary trees* T_1 *and* T_2 *each with* n *interior nodes , the restricted rotation distance* $d_{RR}(T_1, T_2) \leq 12n$.

Proof: Consider the element x of F given by the tree pair diagram (T_1, T_2). We know that x has length less than or equal to $12N(x)$ with respect to the generating set $\{x_0, x_1\}$ and can be thus represented by a string of no more than $12N(x)$ x_0's and x_1's and their inverses. Since these generators correspond to left and right rotations at the root and right child of the root, we have the result.

We also can apply the metric on Thompson's group to get 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 by understanding the notion of a reduced tree pair diagram.

Definition 4.1 *A tree pair diagram* (T_1, T_2) *is* unreduced *if there is an* i *such that the* i*-th and* $i+1$*-th leaves of* T_1 *are both children of the same node and the corresponding* i*-th and* $i+1$*-th leaves of* T_2 *are also both children of the same node. A tree pair diagram which is not unreduced is* reduced.

In the reduction illustrated in Figure 3, the leaves numbered 1 and 2 of T_1 are both children of the same node, as are the corresponding leaves 1 and 2 in T_2, so there is a possible reduction there and also at leaves 5 and 6. After reduction, there are two fewer leaves and the leaves can be renumbered from left to right. So after reduction and relabelling, the tree pair diagram (T_1, T_2) becomes (T_1', T_2'). Note that the rotations that transform T_1 to T_2 are exactly the same rotations needed to transform T_1' to T_2', as the reducible pieces of T_1 and T_2 are carried along unchanged. Those reduced pieces hang from leaves in T_1' and T_2' which, though they may be moved around by the transformation, are not descended into and thus the reducible pieces are carried around unchanged. Applying the lower inequality of Theorem 2, we have the the following, which gives a lower bound on the number of rotations at the root and right child of the root needed to transform one tree to another:

Theorem 4 *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* $d_{RR}(T_1, T_2) \geq \frac{n-1}{3}$.

5 Algorithm for Computing Restricted Rotation Distance

We now describe an efficient algorithm for giving a sequence of rotations to transform T_1 to T_2, rooted

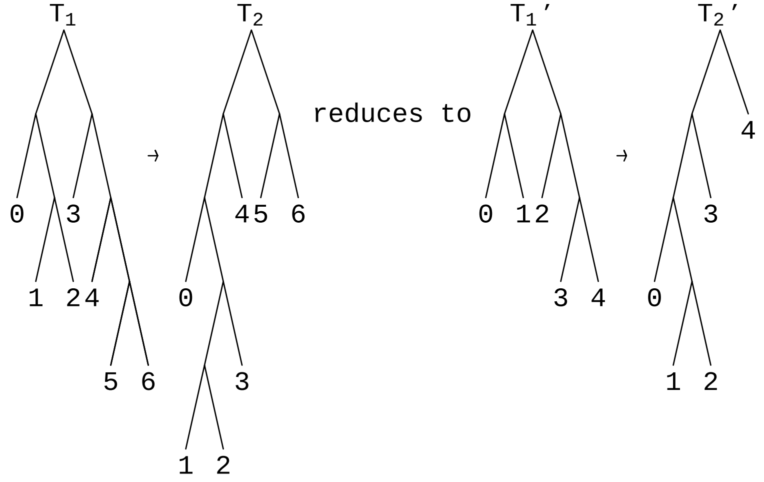

Fig. 3. Example of reduction and relabeling

binary trees with the same number of leaves n. Our rotations are restricted to those at the root and right child of the root. Given T_1 and T_2, we consider the element x in F whose tree pair diagram is (T_1, T_2). We construct the normal form in F by the following process, described in [11]: number the leaves of T_1 and T_2 from left to right beginning at 0. For the ith leaf, we count the maximal length path of left edges beginning at the leaf which does not reach the right side of the tree. That is, we consider the set of ancestors of the ith leaf, and count the number of ancestors which are connected to leaf i by a path consisting entirely of left edges, subtracting one if the most distant such ancestor is on the right side of the tree. Call these lengths r_i for T_1 and s_i for T_2; many of these lengths may be zero. The word $x_0^{r_0}x_1^{r_1}\dots x_n^{r_n}x_n^{-s_n}\dots x_1^{-s_1}x_0^{-s_0}$ represents x in the normal form of the infinite generating set for F. Many of the exponents are possibly zero and introduce nothing into the word. Then, we use the relationship $x_k = x_0^{-k+1}x_1x_0^{k-1}$ to express the generators x_k for $k \geq 2$ in terms of the two generators x_0 and x_1. We then perform the obvious cancellations, canceling $x_0x_0^{-1}$ and $x_0^{-1}x_0$ which can be quickly eliminated in a single pass though the word. (The potential cascading of cancellations such as that encountered in a subword of the form $x_0^2x_1^{-1}x_0x_0^{-1}x_1x_0^{-2}$ cannot occur because of the original form of the word x is in a standard form for the infinite presentation–guaranteeing that all the x_1's precede all the x_1^{-1}'s and that there is a net non-zero exponent for x_0 between the last x_1 and the first x_1^{-1}.) This gives x expressed as a product of x_0's and x_1's which will transform the tree T_2 to the tree T_1. If we want to transform T_1 to T_2, we take the inverse x^{-1} which gives a sequence of rotations transforming T_1 to T_2. This may not be the optimal representation of x realizing $d_{RR}(T_1, T_2)$, but a consequence of Theorem 1 is that this estimate is within a factor of three of the optimum. The steps required to find this expression of rotations between two trees each with n leaves involves: a single traversal of each tree to compute the exponents for each leaf and thus construct the word x, a single pass through the word x to express it in terms of x_0's and x_1's, a single pass through the resulting word to cancel occurrences of $x_0x_0^{-1}$ and $x_0^{-1}x_0$, and a single pass through the resulting word to reverse it. Since each pass requires at most a linear number of steps in n, the algorithm is linear in the size of the given trees.

As an example, given the trees T_1 and T_2 in Figure 4, we estimate the reduced rotation distance and find the sequence of rotations realizing that estimate. First, we find the exponents for T_1. In T_1, leaf 0 has no path of left edges which do not reach the right hand side of the tree (there is a single left edge beginning at leaf 0, but it reaches the right side of the tree). Leaf 1 in T_1 has a single left edge not reaching the right side of the tree, so its exponent is 1. Continuing, we get exponents for T_1 being $(0, 1, 0, 0, 0, 0)$. For T_2, we get exponents of the leaves as $(1, 0, 1, 1, 0, 0)$. Thus, the word x in F representing (T_1, T_2) is $x_0^1 x_1^0 x_2^1 x_3^1 x_4^0 x_5^0 x_5^{-0} x_4^{-0} x_3^{-0} x_2^{-0} x_1^{-1} x_0^{-0}$ which, after omitting terms with exponent 0, is $x_0x_2x_3x_1^{-1}$. We act on the relationship $(T_1, T_2) = x_0x_2x_3x_1^{-1}$ to get $(T_1x_1x_3^{-1}x_2^{-1}x_0^{-1}, T_2) = id$, a representative of the identity transformation and thus two copies of the same tree. Then we use the relationships $x_2^{-1} = x_0^{-1}x_1^{-1}x_0$ and $x_3^{-1} = x_0^{-2}x_1^{-1}x_0^2$ to rewrite this as a sequence of rotations: $x_1x_0^{-2}x_1^{-1}x_0^2x_0^{-1}x_1x_0x_0^{-1}$. We cancel out the occurrences of $x_0x_0^{-1}$ to get $x_1x_0^{-2}x_1^{-1}x_0x_1^{-1}$. That sequence of 6 rotations transforms T_1 to T_2 as shown. For this small example, that turns out to be the unique minimal sequence of rotations and $d_{RR}(T_1, T_2) = 6$, but in general, there may be many such minimal sequences of rotations and the algorithm may find a longer sequence than one of the minimal ones. A better understanding of the metric properties of the group F might lead to algorithms to better compute the restricted rotation distance.

The author would like to thank Joel Hass and Katherine St. John for helpful conversations and the anonymous referees for helpful suggestions.

References

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

[2] K. Culik II, D. Wood, A note on some tree similarity measures, Inform. Process. Lett. 15 (1) (1982) 39–42.

[3] J. Pallo, On the rotation distance in the lattice of binary trees, Inform. Process. Lett. 25 (6) (1987) 369–373.

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

[5] E. Mäkinen, On the rotation distance of binary trees, Inform.

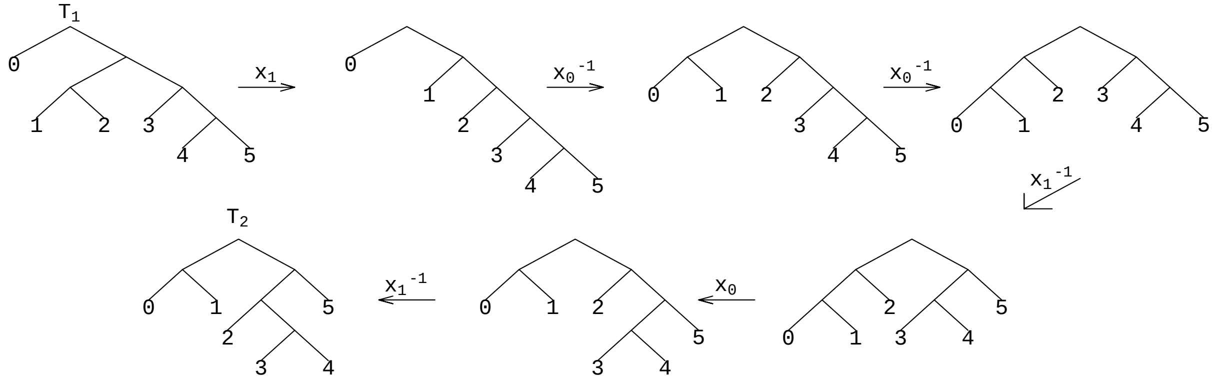

Fig. 4. Sequence of rotations to change T_1 to T_2

Process. Lett. 26 (5) (1988) 271–272.

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

[7] F. Luccio, L. Pagli, On the upper bound on the rotation distance of binary trees, Inform. Process. Lett. 31 (2) (1989) 57–60.

[8] S. Hanke, T. Ottmann, S. Schuierer, The edge-flipping distance of triangulations, J.UCS 2 (8) (1996) 570–579.

[9] 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.

[10] A. Bonnin, J. Pallo, A shortest path metric on unlabeled binary trees, Pattern Recognition Letters 13 (1992) 411–415.

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

[12] J. Burillo, Quasi-isometrically embedded subgroups of Thompson's group F, J. Algebra 212 (1) (1999) 65–78.

[13] 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.

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

[15] K. 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), Springer, New York, 1992, pp. 137–163.

[16] M. Gromov, Hyperbolic groups, in: S. M. Gersten (Ed.), Essays in Group Theory, Springer, 1987.

[17] 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.