Journal of Graph Algorithms and Applications
http://jgaa.info/ vol. 27, no. 1, pp. 19–33 (2023)
DOI: 10.7155/jgaa.00611

Restricted rotation distance between k-ary trees

Sean Cleary

The City College of New York and The CUNY Graduate Center

Submitted:	June 2022	Reviewed:	July 2022	Revised:	December 2022
Accepted:	February 2023	Final:	February 2023	Published:	February 2023
Article type:	Regular paper	Communicated by:	G. Liotta		

Abstract. We study restricted rotation distance between ternary and higher-valence trees using approaches based upon generalizations of Thompson's group F. We obtain bounds and a method for computing these distances exactly in linear time, as well as a linear-time algorithm for computing rotations needed to realize these distances. Unlike the binary case, the higher-valence notions of rotation distance do not give Tamari lattices, so there are fewer tools for analysis in the higher-valence settings. Higher-valence trees arise in a range of database and filesystem applications where balance is important for efficient performance.

1 Introduction

Rotation distance for binary trees measures the difference in tree shape in terms of a number of fundamental small changes. Changing tree shape via small steps has important role in arranging unbalanced tree to balanced trees, which are needed for optimal performance of search algorithms. A number of questions about rotation distance have been studied. No polynomial-time algorithm is known to compute rotation distance exactly, though there are polynomial-time estimation algorithms of Pallo [32], Rogers [37], Baril and Pallo [1], Cleary and St. John [18], and Cleary and St. John [17] have shown that the rotation distance problem is fixed-parameter tractable. The question of bounding rotation distance between two trees has been addressed by a number of researchers, including Sleator, Tarjan and Thurston [38], who used hyperbolic volume methods to show that no more than $2n-6$ rotations are needed to transform one rooted binary tree with n nodes into any other and that this bound is achieved for all sufficiently very large values of n and thus is optimal in this asymptotic sense. Pournin [35] showed that the bound is optimal in fact for $n \geq 11$.

There are natural analogues of rotation distance for higher-valence trees but the techniques for analysis which are effective for binary trees do not work well in the more general setting.

Support for this project was provided by a PSC-CUNY Award, jointly funded by The Professional Staff Congress and The City University of New York.

E-mail address: scleary@ccny.cuny.edu (Sean Cleary)

Particularly, as described by Pallo [33], the set of rotation transformations for binary trees gives rise to a Tamari lattice on binary trees, but for any higher valence trees, the poset arising from a natural generalization of the rotation transformations is not a Tamari lattice, as described by Sunik [40]. Huguet [28] showed that the analogous posets in the higher valence setting are not even lattices, so many of the techniques for the binary case are ineffective. Pallo [33] shows that the posets arising from some rotation transformations for n-ary trees do satisfy the weaker condition of being a χ-lattice in the sense of Leutola and Nieminen [30] which leads to some algorithms for computing meets and joins there, but the analysis is quite difficult and the understanding is far less complete than in the binary cases. Even in the ternary case, very little progress has been made using these approaches.

Restricted rotation distance is similar to rotation distance in that it measures the difference in rooted binary tree shape by counting the minimum number of rotations needed to transform one tree to another, but it differs from ordinary rotation distance in that instead of rotations being permitted anywhere in the tree, rotations are permitted only at two distinguished nodes– the root node and the right child of the root (if present). Restricted rotation distance is in general significantly larger than rotation distance, but there are a number of similarities and differences, as described by Cleary [11] and Cleary and Taback [19, 20]. The metric properties of the abstract group known as Thompson's group F with respect to its standard finite generating set give an effective means of studying restricted rotation distance. Restricted rotation distance can be computed exactly in linear time using ideas developed by Fordham [25], and an exact sequence of rotations needed to realize this distance can also be found with a linear-time algorithm. For two trees with n interior nodes, the ordinary rotation distance is sharply bounded by $2n-6$ in an asymptotic sense, as shown by Sleator, Tarjan and Thurston [38]. For restricted rotation distance between two trees with n nodes, Cleary and Taback show that the restricted rotation distance is bounded by $4n-8$ and that this bound is sharp and achieved for $n \geq 3$. Cleary and Taback [20] and Pallo [34] also study a number of other restrictions on possible locations where rotations are permitted with various notions of "right-arm" and "restricted right-arm" rotation distance, and Luccio, Mesa Enriquez and Pagli [31] study "level restricted" rotation distances, where rotations are only permitted at the top levels of a tree. Cleary, Luccio, and Pagli [13] give sharper bounds in specific cases of restricted right-arm rotation distances.

We mention that the well-developed connections between edge-flip distance between triangulations of polygons and rotations of binary trees (see [38]) do not have well-studied analogs in higher valence cases. Binary trees correspond to subdivision of polygons into triangles, and higher-valence trees can correspond to subdivisions of polygons into larger subpolygons, such as quadrilaterals in the ternary case and pentagons in the quarternary case, which have not drawn the extensive study that subdivisions into triangulations have.

Here, we consider questions about the rotation distance between higher valence trees. Many methods which work effectively for binary trees fail for higher valence trees, as described above, but the approaches using Thompson's group F can be adapted to effectively understand some of the properties of appropriate notions of restricted rotation distance between higher valence (regular k-ary) trees by considering generalizations of Thompson's group F, known as $F(k)$. The group $F(2) = F$ is the usual well-studied Thompson's group F. In this interpretation of F, rooted binary trees correspond to dyadic subdivisions of the unit interval, and group elements correspond to piecewise-linear interpolations between dyadic subdivisions. The resulting slopes are necessarily powers of 2 and the places where slope changes must be dyadic rationals. The corresponding interpretation of $F(3)$ corresponds to ternary trees with triadic subdivisions, for example. The group elements in $F(3)$ are piecewise-linear interpolations between triadic subdivisions, with slopes

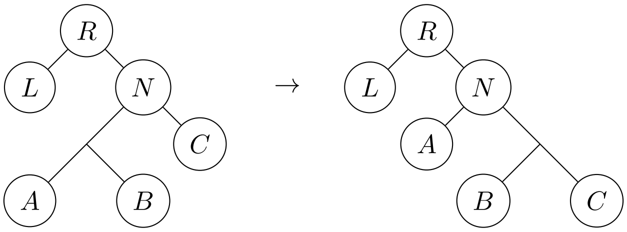

Figure 1: Right rotation at node N in a binary tree transforms the left tree to the right tree, where A, B, C, and L can be leaves or subtrees. Left rotation at N performs the inverse operation taking the right tree to the left one.

being powers of 3 and slopes changing only at triadic rationals $\mathbf{Z}[\frac{1}{3}]$. Trees with higher valence k have corresponding Thompson's groups with slopes in k^i and breakpoint sets contained in the $k-$adics $\mathbf{Z}[\frac{1}{k}]$.

Here, we describe a linear time algorithm to compute restricted rotation distances exactly and give sharp bounds on restricted rotation distances in these higher-valence settings, for rooted trees of all valences k.

2 Background and definitions

A rooted k-ary tree is a rooted, ordered tree where each interior node has exactly k children, which are each either interior or exterior nodes. In the following, by *tree*, we mean a rooted k-ary tree. We call exterior nodes *leaves* and number the leaves in increasing order from left to right starting at leaf 1. The *right side* of the tree is the root and those nodes connected to the root by a path consisting entirely of right edges; similarly we have the *left side* of a tree. A rooted k-ary tree with n interior nodes will have $n(k-1)+1$ leaves.

There are several possible notions of what constitutes a "basic rotation" in a higher-valence tree. In a binary tree, there are only two possible children to consider, and rotation promotes a grandchild from one of those children and demotes the other child to a grandchild. In the binary case, the standard straightforward choice is to allow a left rotation which promotes the right child of the right child of node N to be the new right child of N, moves the left child of the right child of N to be the right child of the left child of N, and demotes the left child of N to be the left child of the left child of N, as pictured in Figure 1. If we are permitted to operate at any node present in the tree, and allow both right and left rotations when possible, we can transform any rooted binary tree to any other rooted binary tree with the same number of nodes. The standard *rotation distance* between two binary trees with the same number of nodes is the minimum number of rotations needed to transform T_1 to T_2, where the rotations can be performed at any node.

For higher valence trees, since there are more children, there are a number of differing possible notions for what constitutes a corresponding "basic rotation." It is clear that in order to be able to transform a given tree to any other tree with the same number of leaves, it will be necessary to allow rotations which collectively can promote and demote into each of the k children of a node. Unlike in the binary case, there are several competing reasonable choices for what rotations to allow during these transformations. We could consider promotion and demotion, in order-preserving manners, from any of combination of two of the children, giving $\binom{k}{2}$ possible "basic rotations" at a node.

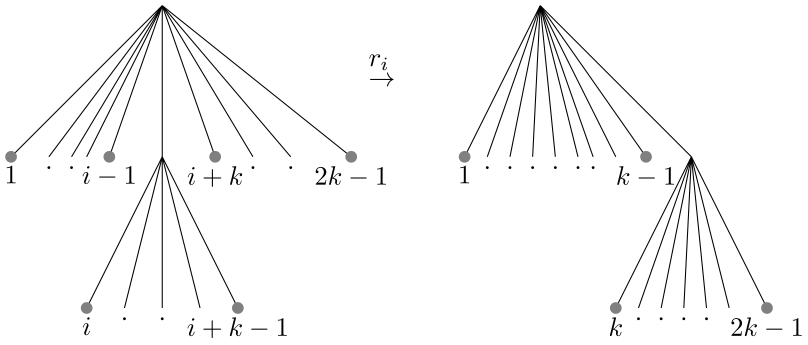

Figure 2: The $k-1$ fundamental right rotations r_i at the root in a tree of valence k, for $1 \leq i \leq k-1$ transform the left tree to the right tree. The fundamental left rotations are the inverse operations.

Or we could allow promotion and demotion (again, preserving order) from combinations which include the rightmost child and any of the other children of the node, giving $k-1$ possible "basic rotations". Or, as Pallo does in the ternary case detailed in [33], we could only allow rotations which promote from and demote into adjacent children, also giving $k-1$ possible "basic rotations." Each of these alternatives for fundamental allowed rotations will give rise to a different notion of rotation distance on k-ary trees, as they each give a different set of fundamental moves which count as one step in the associated metrics. The role of the right side of the tree and of the "all-right tree" in the work of Culik and Wood [22] on binary trees suggest that we allow order-preserving promotions and demotions involving various children and the rightmost child into any of the other children of the node. This choice will allow the best use of the connections with the generalizations of Thompson's group F. Thus, our choice of "basic rotations" is to allow transformations from the left-hand tree in Figure 2 to any of the right-hand trees and also to allow the inverses ("left rotations") of those transformations. And thus, we define the *rotation distance* $d_R(T_1, T_2)$ between two k-ary 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 present and are of the types described in Figure 2.

As described in [11] and [19], it is not necessary to allow rotations at all possible nodes to be able to transform any given tree to any other tree with the same number of nodes. Here, we consider how restricting to a smaller set of allowed rotations leads to the notion of restricted rotation distance.

As pointed out above, to ensure that we allow sufficiently many rotations to transform any tree to any other with the same number of nodes, we will need at least $k-1$ rotations (and their inverses) to collectively promote and demote into any of the k children. But those rotations alone at the root are insufficient, as will be shown below. So we must allow at least one additional rotation. If we allow merely a single additional rotation at the right child of the root, shown in Figure 3, we will be able to transform any tree to any other, as will be shown below. Further, any set of fundamental rotations with fewer than k rotations will be insufficient for the desired goal of being able to transform any tree of size n into any other tree of size n. There are other possible minimal sets of fundamental rotations but the choice of these particular ones will allow the efficient computation described in Section 5. Thus, we allow the $k-1$ different rotations pictured in Figure 2 at the root, together with the single rotation at the right child of the root from Figure 3 as the basic rotations in the k-ary case. So we define *restricted rotation distance* $d_{RR}(T_1, T_2)$ between two

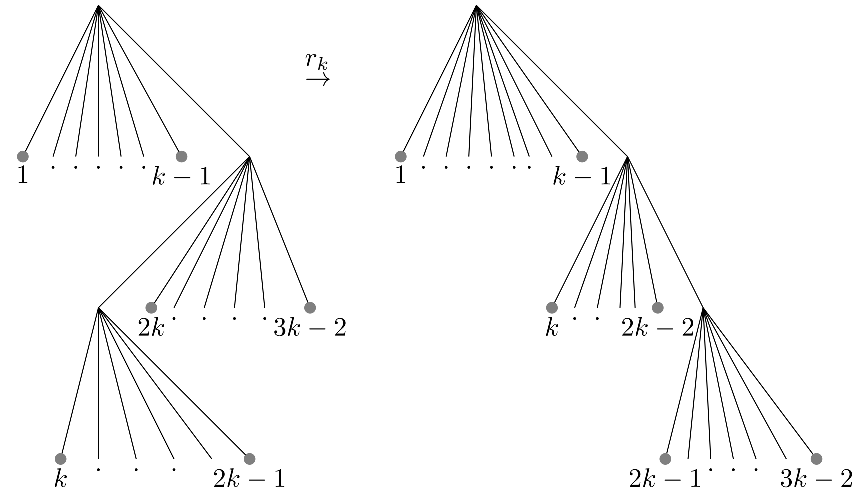

Figure 3: A preferred single right rotation r_k at the rightmost child of the root in a tree of valence k.

rooted k-ary 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 allowed are from the set of k possible rotations just described.

As in the binary case, 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 (the root and the right child of the root) to accomplish the equivalent operation of a single rotation at a single node at a great distance from the root, via a conjugation process. Though the number of steps required to perform the transformation is higher, the individual steps are very simple and only of k possible types. To understand restricted rotation distance, we study Thompson's group $F(k)$.

3 Thompson's Groups $F(k)$ and Tree Pair Diagrams

Just as Thompson's group $F = F(2)$ is useful for understanding restricted rotation distance in the binary case with 2 children per node, the generalized Thompson's groups $F(k)$ with $k \geq 3$ are useful for analyzing restricted rotation distances in the higher valence case with k children per node. Thompson's groups $F(k)$ can be defined in three equivalent ways, all of which are useful in understanding the structure of the group.

Analytically, we can define $F(k)$ as a group of piecewise-linear homeomorphisms of the unit interval. The group $F(k)$ is 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 k-adic rationals $\mathbf{Z}[\frac{1}{k}]$, and, away from the singularities, the slopes are integral powers of k.

Combinatorially, we can define the group $F(k)$ in terms of generators and relations. In terms of generators and relations, $F(k)$ has infinite presentation:

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

There is a set of normal forms for elements of F 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 $r_i, s_i > 0$, $i_1 < i_2 \dots < i_k$ and $j_1 < j_2 \dots < j_l$. This normal form is unique if we further require a reduction condition that when both x_i and x_i^{-1} occur, so does at least one of $x_{i+1}, x_{i+1}^{-1}, x_{i+2}, x_{i+2}^{-1}, \dots x_{i+k-1}, x_{i+k-1}^{-1}$ appears, as discussed by Burillo, Cleary and Stein [6]. The relations provide a quick and efficient manner to rewrite a word into normal form, and form a complete rewriting system, as described by Brown [5] and Stein [39].

There is also a finite presentation for $F(p)$ which is essentially a sub-presentation of the infinite presentation. Since x_0 conjugates x_1 to x_k and similarly all higher-index generators x_l with $l \geq k$ are also conjugates of one of $x_1, \dots x_{k-1}$ by an appropriate power of x_0, the k generators $\{x_0, \dots x_{k-1}\}$ suffice to generate the whole group. Furthermore, though the presentation of $F(p)$ given above is infinite both in terms of the number of generators and relations, $F(p)$ is finitely presentable, requiring only the generators $\{x_0, \dots x_{k-1}\}$ and a fundamental set of $k(k-1)$ relations, as all of the remaining infinitely many relations are consequences of those fundamental ones, as described by Brown [4] and Stein [39].

Geometrically, we can regard $F(k)$ as equivalence classes of pairs of rooted k-ary trees. The elements of $F(k)$ are represented as equivalence classes of tree pair diagrams of valence k. A *tree pair diagram* is a pair of k-ary trees with the same number of leaves, as described in [7]. Such a tree can be regarded as a procedure for constructing a subdivision of the unit interval by successive division of intervals into k equal length pieces. A pair of such trees (S, T) gives an element of $F(k)$ which can be regarded from the analytical perspective by considering the element f which is the piecewise-linear homeomorphism which realizes the linear 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 detail in the binary case by Cannon, Floyd and Parry [7] and Cleary [12].

These generators correspond to the fundamental set of basic rotations chosen for restricted rotation distance, where x_i corresponds to the rotation r_{i+1} shown in Figure 2 or 3 owing to the choices of numbering leaves starting from 0 (commonly done in the group-theoretic descriptions of $F(k)$) or from 1 as done more commonly in data structures and as principally done here. The additional generators x_n for $n \geq k$, which occur in the infinite generating set but not the standard finite generating set, correspond to right rotations at nodes further down on the right-hand side of the tree. Specifically, the generator x_n rotates rightward from the child $n \mod (k-1)$ from the left of the rightmost node on the level $\lfloor \frac{n}{k-1} \rfloor$ to the rightmost child of the rightmost node at that level.

As in the binary case, there is a natural notion of reduction and a natural group operation. A tree pair diagram (T_1, T_2) is *unreduced* if there is an i such that the i-th through $i+k-1$th leaves of T_1 are all children of the same node and the corresponding i-th through $i+k-1$th leaves of T_2 are also children of a single node in T_2. There is a natural reduction process by which the unreduced tree pair (T_1, T_2) is replaced with a tree pair (T_1', T_2') where we obtain the primed trees by deleting the leaves i-th through $i+k-1$th leaves, replacing the parent node of the deleted leaves with a leaf, and renumbering. A *reduced* tree pair diagram is not unreduced. For the equivalence relation where two tree pairs (S, T) and (U, V) are considered equivalent if they have a common tree which can be obtained via a sequence of reductions, there is a unique reduced tree pair diagram in each equivalence class, just as there is in the binary case.

To obtain a group, we define composition analogously to the binary case. The group operation is composition; it may be necessary to expand trees (by a succession of moves which change leaves to interior nodes and add corresponding new leaves in both trees in the pair, making intermediate unreduced diagrams) to create unreduced representatives of elements of $F(p)$ to perform composition (see [7, 12] for details in the binary case.)

To convert between the geometric description of elements of $F(p)$ to the combinatorial one, we use leaf exponents. The *exponent* $E(i)$ of the ith leaf is the integral length of the longest path consisting entirely of left edges beginning at the leaf which does not reach the right side of the tree. To construct the normal form of an element given in terms of a k-ary tree pair diagram in terms of the infinite generating set for $F(k)$ we use a similar process to that described in [7] and [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. We call these exponents s_i for T_1 and r_i for T_2. Many of these exponents may be zero as the only way to have a non-zero exponent is if the corresponding leaf is the leftmost leaf of its parent node and for its parent node is not on the right side of the tree. 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 $F(k)$.

Given any pair of rooted $k-$ary trees (T_1, T_2) with the same number of leaves, we can consider that as representing an element of $F(k)$. And given any element f of $F(k)$, 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. Thus, there is a one-to-one correspondence between reduced tree pair diagrams and elements of $F(k)$.

4 The Metric on $F(k)$ and Restricted Rotation Distance

Given a group G with a finite generating set, we define the *length* of a word g in G as minimum length over all possible representatives of g as words in terms of the generators and their inverses. 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. Geometric group theory is concerned with understanding the metric properties of groups and the consequences of metric hypotheses on groups; see Clay and Margalit [10] or Epstein, Cannon, Holt, Levy, Paterson and Thurston [24] for an introduction.

Burillo, Cleary and Stein [6] estimated the word metric on $F(k)$, showing that word length of a group element is proportional to the number of nodes in the reduced tree pair diagram representing the element. Fordham developed a method, described in [26], for computing the word length in $F(k)$ exactly, using a method based upon his method for computing word length in F exactly [25]. We use his method here to approach restricted rotation distances.

First, we note that no set of less than k fundamental rotations can suffice to give all possible transformations between rooted k-ary trees of the same size.

Proposition 1 *Any set of rotations of size smaller than k cannot be sufficient to transform one given tree with n nodes into all other possible trees of n nodes.*

Proof: This follows from the fact that the group $F(k)$ is a k-generator group with abelianization $\mathbb{Z}^k$, see Brown [4]. Given a set of less than k rotations at specified nodes, the subgroup of $F(k)$ generated by those rotations interpreted as group elements cannot be the whole group, as the rank of the abelianization will be too small, and thus there are elements not in the subgroup. The elements not in the subgroup give tree pairs which are not expressible as products of the alleged generating set and thus give pairs of trees which are not able to be transformed via those permitted rotations. ☐

Second, we note that our collection of k rotations denoted pictured in Figures 2 and 3 form a sufficient set of rotations to perform any desired transformation.

Proposition 2 *Given two trees of the same size T_1 and T_2, there is a sequence of rotations of the types r_i with $1 \leq i \leq k$ which transform T_1 into T_2.*

Proof: The tree pair (T_1, T_2) gives an element w of Thompson's group $F(k)$. In $F(k)$, the word w can be represented by a sequence of generators $g_1 g_2 \dots g_N$, each of which corresponds to a rotation of the permitted type. Applying the N corresponding rotations to T_1 in order will yield T_2. ☐

The connection between word length in Thompson's groups comes from regarding two k-ary trees T_1 and T_2 as an element $f = (T_1, T_2)$ in $F(k)$ and the following:

Proposition 3 *The restricted rotation distance between two k-ary trees T_1 and T_2 is exactly the word length with respect to the standard generators in $F(k)$ of the group element f represented by the tree pair (T_1, T_2).*

Proof: By expressing f as a product of the generators of length m then that expression for f in terms of the generators gives that change broken down into m rotations at the root and right child of the root, of the appropriate allowed types. One potential complication is that generators for elements in $F(k)$ may be applied to trees which are too small to accomplish the rotation- in the case of $F(k)$, we merely choose an appropriately larger unreduced representative of the element and apply the generator to that. This increases the number of nodes in the tree pair diagram. There is also a natural notion of reduction, which can result in the decrease of the number of nodes during the application of a sequence of generators. For rotation distance, however, the number of nodes must remain constant during the transformation. But by Theorem 2.5 of [26], for a geodesic path representing $f = (T_1, T_2)$, the number of nodes in the sequence of reduced representatives does not increase and we may opt not to perform any reductions, so the tree sizes will indeed remain constant during the sequence if we use this approach. Thus, the restricted rotation distance between T_1 and T_2 will be no more than m. Conversely, were the restricted rotation distance between T_1 and T_2 to be less than m, we could use the applicable rotations to construct a shorter representative of the group element (T_1, T_2) in $F(k)$, contradicting the minimality of word length there. So indeed the word length and the restricted rotation distance coincide. ☐

5 Fordham's calculation of word length in $F(p)$

Fordham [25] presented an algorithm for finding the exact length of an element of F, later generalized to find the exact length of an element of $F(p+1)$, described in [26]. This process can give a linear-time algorithm for finding rotation distance between k-ary trees, for arbitrary fixed k. Here we describe the process and show that the running time is linear in the size of the trees.

Theorem 1 *There is a linear-time algorithm to compute the restricted rotation distance between two k-ary trees T_1 and T_2 each with n internal nodes.*

Proof: First, we note that the interior nodes of a tree T can be given orders in a manner somewhat analogous to the natural infix order where the nodes in the left subtree of a node n precede it and the nodes in the right subtree of n follow it. In the binary case, there was exactly one left child and one right child. In the higher valence case, we segregate the children of a node into "left children" and "right children." Each node will have at least one child of each type but the particular allocation of the number of left and right children will depend upon the node type in a way related to its type relative to residues modulo $k-1$. We give the following recursive definitions of node type, where the node types broadly are root (considered as a distinguished node of type L), left nodes L, right nodes R, and $k-1$ types of middle nodes of types $M^1, \ldots, M^{k-1}$. (Note that [26] uses valence $p+1$ trees, and thus resulting in a different notation convention for node types.)

The node type of the root node is 'root', considered as type L and it has one left child of type L and $k-1$ right children of types $M^1, \ldots, M^{k-2}$ and R.

A node of type L has one left child of type L and $k-1$ right children of types $M^1, \ldots, M^{k-1}$.

A node of type R has one left child of type M^{k-1} and $k-1$ right children of types $M^1, \ldots, M^{k-2}$ and R.

A node of type M^i has $k-i$ left children of types $M^i \ldots M^{k-1}$ and i right children of types $M^1, \ldots, M^i$.

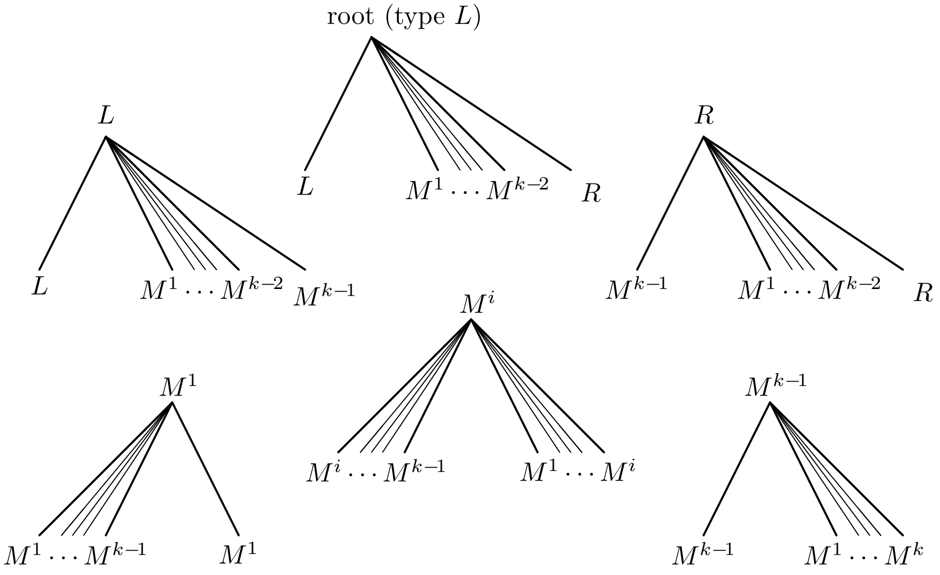

Figure 4: Node types and corresponding child nodes. Though the number of total children is fixed at k for all nodes, note that not all node types have the same number of left and right children, as indicated by the spacing above.

This recursive definition of node types is described in Figure 4.

A tree pair (S, T) is *reduced* if there is no sequence of leaves $i+1, i+2, \ldots, i+k$ which are leaves of a single node in S and also leaves of a single corresponding node in T. If a tree pair is not reduced, there is a unique reduced tree pair equivalent under a natural reduction process, as described in [11] in the binary case. Further, if an unreduced tree pair (S', T') reduces to a tree pair (S, T), the rotations required to transform S' to T' are exactly the same as those required to transform S to T, and all nodes where rotations are actually performed will be present in

		Target Node Type					
		L_{L}	$R_\emptyset$	R_{R}	R_j	$M^i_\emptyset$	M^i_j
Source node type	L_{L}	2	1	1	1	2	2
	$R_\emptyset$	1	0	2	2	1	3
	R_{R}	1	2	2	2	1	3
	R_{j_1}	1	2	2	2	3	3
	$M^{i_1}_\emptyset$	2	1	1	1	2	2
	$M^{i_1}_{j_1}$	2	3	3	3	4	4
	R_{j_2}	1	2	2	2	1	3
	$M^{i_2}_\emptyset$	2	1	1	3	2	4
	$M^{i_2}_{j_2}$	2	3	3	3	2	4

Table 1: Weights for node pairs with $j_1 \leq i < j_2$, $i_1 < j \leq i_2$.

both pairs at the appropriate times, using an argument parallel to that in [11]. The additional reducible nodes and leaves are merely carried along unchanged during the transformation. Note that reducible nodes are considered to not exist for the purposes of classify node types more finely where the notion of the type of successor node is involved, just as in the binary case.

There is a natural inorder traversal of the nodes of these types of tree where all of the left children are visited before the node, and all of the right children visited afterward, in the natural left-to-right order when there are multiple left and right children occurring. By "successor node" in the following, we mean the immediate successor of a node in the resulting inorder traversal of the tree, ignoring leaves.

We further subdivide the node types after reduction as follows.

The leftmost node is of type L_0. Left nodes with a preceding node are of type L_L. This will include all left nodes except a single left node of type L_0. Right nodes all of whose non-reducible successors are right nodes are of type R_0. Right nodes whose immediate non-reducible successor is a right node but which have at least once middle node nonreducible successor are of type R_R. Other right nodes are of type R_j, where the leftmost nonreducible successor of the node is of type M^j.

Middle nodes of type M^i are classified further into finer types as follows. Middle nodes are of type $M^i_\emptyset$ if they have no nonreducible successor nodes, or of type M^i_j if the leftmost nonreducible successor node is of type M^j. For M^i_j, it follows from the node type definitions that $j \geq i$.

To measure the restricted rotation distance between two input trees S' and T' with n interior nodes each, we use the following process:

1. Form the reduced tree pair (S, T) via a sequence of reductions if needed, which can be done in a linear-time preprocessing step resulting in two trees of size l, with $l \leq n$.

2. Determine the node types of the internal nodes in S and T, which can be done efficiently via an inorder traversal of the tree S and then one for the tree T, done in time linear in l, and constructing two lists of node types $(S_0, S_1 \ldots, S_l)$ and $(T_0, T_1, \ldots, T_l)$, ordered corresponding to the inorder traversal described above.

3. Traverse the two lists and for each node pair type in the list, add to the restricted rotation distance the entry from the appropriate matrix of node pair weights given in Table 1 again time linear in l.

The cumulative sum of all node pair weights will give the restricted rotation distance between the original trees S' and T', just as the corresponding process gives restricted rotation distance in the binary case. □

We also note that not only is the restricted rotation distance efficiently findable, but further a minimal sequence of permitted rotations is also obtainable quickly.

Theorem 2 *There is an algorithm linear in the number of nodes n of two valence k trees T_1 and T_2 to find a sequence of permitted rotations to realize the restricted rotation distance from T_1 to T_2.*

Proof: Just as in the binary case as described in [11] and [19], the type of pairing changes between the applications of the generators is bounded in complexity as a constant obtained by multiplying the size of the table by the number of allowed rotations, together with some cases from classifying the presence or absence of particular nodes. Thus, we can construct larger versions of Table 1 which will give a reduction of total weight one in each case, resulting in a minimal length path, just as in the binary case described in [25]. □

We note that the related problem of enumerating all possible minimal length paths is not feasible in linear time, as there are specific examples where the total number of possible such paths grows at least quadratically. Work of Elder, Fury, and Rechnitzer [23] shows that the number of geodesics in Thompson's group F (corresponding to the binary case) grows exponentially and at a rate at least as large as that of elements. For the number of geodesics, the best known algorithms [23] run in exponential time and presently only the first 22 terms of the geodesic growth sequence are known. The computational evidence suggests that the exponential growth rate of the number of geodesics is measurably larger than that for group elements. If so, it means that there are necessarily examples with large numbers of geodesics and resultantly, tree pairs with many possible minimal length paths realizing the restricted rotation distances. We anticipate that the higher valence cases have similar behavior, where there being tree pairs have many possible minimal length paths realizing restricted rotation distances.

6 Applications and related questions

Having a linear-time algorithm for restricted rotation distance in the ternary case opens up analyses parallel to that for restricted rotation distance in the binary case. Some of the generalizations are immediate. For example, the sharp bounds of Cleary and Taback [19] for binary restricted rotation distance are precisely the same in all of the higher-valence instances:

Theorem 3 *Given two rooted k-ary trees T_1 and T_2 each with $n \geq 3$ internal nodes, the restricted rotation distance between them satisfies $d_{RR}(T_1, T_2) \leq 4n - 8$. Furthermore, if (T_1, T_2) form a reduced pair, then $n - 2 \leq d_{RR}(T_1, T_2) \leq 4n - 8$ with these bounds being realized for all $n \geq 3$.*

Proof: The arguments are parallel those of [19] relating to possible reductions and analysis of pair weight changes. The values in the table of weights for the k-ary case range from 0 to 4, the same as in the binary case. There will always be one node pair of weight 0 corresponding to the first node pair which is necessarily of type $(L_\emptyset, L_\emptyset)$ and there can be at most one other node pair of weight 0 in a reduced tree pair diagram, a final node pair of type $(R_\emptyset, R_\emptyset)$. This gives the lower and upper bounds, and the examples of Cleary and Taback [19] in the binary case apply here to realize these bounds for all $n \geq 3$. □

There are a number of questions arising related to rotations and balance for higher-valence trees. Just as rooted binary trees play a foundational role in search trees, higher-valence trees arise in a range of settings related to information storage for databases and filesystems. The applications of higher valence trees include tries, which play an important role in predictive text applications, see Bentley and Sedgewick [2, 3]. The structure of 2-3 trees (see Knuth [29]) can be used to store data efficiently, and these trees have either two or three child nodes at each internal node. B-trees have valence up to a fixed m and arise in database applications and disk storage applications, see Comer [21], and Rodeh, Bacikm and Mason [36] describing the use of B-trees in Linux filesystems. B+ trees often have ordered lists of 100 child nodes or more to reduce the number of levels of indexing for large databases, see Giampaolo [27]. All of these higher-valence trees require reasonable balance to guarantee efficient worst-case search time complexity, and so questions of balance arise in each of these higher-valence settings.

With the linear-time methods described here, many of the questions which have been addressed for binary trees with respect to restricted rotation distance have natural generalizations to higher valence cases. For example, the investigations of Chu and Cleary [8, 9] and Cleary, Passaro, and Toruno [15] about expected numbers of conflicts and maximum vertex valence in rooted triangulations, the understanding of Cleary and Nadeem [14] about the distribution of restricted rotation distances, as well as the associated results of Cleary, Elder, Rechnitzer and Wong [16] on the asymptotics of such conflicts and partitions, were all developed for the binary case. These questions are open in the higher-valence cases, but it is not clear which may be readily approachable with the tools described here. The connections between rooted trees and the dual triangulations of regular polygons which are useful in several of these works do have generalizations in the higher tree-valence cases to subdivisions of regular polygons into quadrilaterals, pentagons, and so on, but the combinatorial arguments are more complicated and not as well-understood.

Though presumably some of these phenomena have straightforward generalizations to the higher valence cases, that is by no means assured and in fact the Thompson's group versions $F(p)$ have a number of more complicated pathological behaviors, both cohomological and metric, not seen in the much-better studied Thompson's group F.

Many of these questions are wide open in settings where there are trees of variable valence. In the fixed-valence setting, the tools connected to Thompson's groups can be brought to bear, but in the variable-valence cases, these approaches do not show nearly as much promise. The Thompson's group approaches rely upon the symmetry and regularity of the trees. The only groups in this family for which there are exact efficient algorithms for the word metric require constant fixed valences. There are natural group structures associated to the cases where the the node valences can change such as Thompson's group $F(2,3)$ (see Stein [39] for background and Wladis [41, 42, 43] for understanding some of the additional complexities of the word metric in these groups $F(n,m)$ and $F(n_1, n_2, \ldots, n_m)$ which are significantly more complicated), the associated rotation operations do not correspond as neatly with a natural notion of rotation in variable-valence trees. Furthermore, in these settings, the understanding of the word metric is far from complete. Thus, it seems less likely that Thompson's group based approaches are likely to have great success in the variable-valence settings and different approaches may be more effective.

References

[1] J.-L. Baril and J.-M. Pallo. Efficient lower and upper bounds of the diagonal-flip distance between triangulations. *Inform. Process. Lett.*, 100(4):131–136, 2006. `doi:10.1016/j.ipl.2006.07.001`.

[2] J. L. Bentley and R. Sedgewick. Fast algorithms for sorting and searching strings. In *Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (New Orleans, LA, 1997)*, pages 360–369. ACM, New York, 1997.

[3] J. L. Bentley and R. Sedgewick. Ternary search trees. *Dr. Dobbs Journal*, 1998.

[4] K. S. Brown. Finiteness properties of groups. In *Proceedings of the Northwestern conference on cohomology of groups (Evanston, Ill., 1985)*, volume 44, pages 45–75, 1987. doi:10.1016/0022-4049(87)90015-6.

[5] 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)*, volume 23 of *Math. Sci. Res. Inst. Publ.*, pages 137–163. Springer, New York, 1992. doi:10.1007/978-1-4613-9730-4_6.

[6] J. Burillo, S. Cleary, and M. I. Stein. Metrics and embeddings of generalizations of Thompson's group *F*. *Trans. Amer. Math. Soc.*, 353(4):1677–1689, 2001. doi:10.1090/S0002-9947-00-02650-7.

[7] 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.

[8] T. Chu and S. Cleary. Expected conflicts in pairs of rooted binary trees. *Involve*, 6(3):323–332, 2013. doi:10.2140/involve.2013.6.323.

[9] T. Chu and S. Cleary. Expected maximum vertex valence in pairs of polygonal triangulations. *Involve*, 8(5):763–770, 2015. doi:10.2140/involve.2015.8.763.

[10] M. Clay and D. Margalit, editors. *Office hours with a geometric group theorist*. Princeton University Press, Princeton, NJ, 2017.

[11] S. Cleary. Restricted rotation distance between binary trees. *Inform. Process. Lett.*, 84(6):333–338, 2002. doi:10.1016/S0020-0190(02)00315-0.

[12] S. Cleary. Thompson's group. In *Office hours with a geometric group theorist*, pages 331–357. Princeton Univ. Press, Princeton, NJ, 2017.

[13] S. Cleary, F. Luccio, and L. Pagli. Refined upper bounds for right-arm rotation distances. *Theoret. Comput. Sci.*, 377(1-3):277–281, 2007. doi:10.1016/j.tcs.2007.02.024.

[14] S. Cleary and H. Nadeem. Distributions of restricted rotation distances. *Art Discrete Appl. Math.*, 5(1):Paper No. 1.03, 16, 2022. doi:10.26493/2590-9770.1374.bf0.

[15] S. Cleary, J. Passaro, and Y. Toruno. Average reductions between random tree pairs. *Involve*, 8(1):63–69, 2015. doi:10.2140/involve.2015.8.63.

[16] S. Cleary, A. Rechnitzer, and T. Wong. Common edges in rooted trees and polygonal triangulations. *Electron. J. Combin.*, 20(1):Paper 39, 22, 2013. doi:10.37236/2541.

[17] S. Cleary and K. St. John. Rotation distance is fixed-parameter tractable. *Inform. Process. Lett.*, 109(16):918–922, 2009. doi:10.1016/j.ipl.2009.04.023.

[18] S. Cleary and K. St. John. A linear-time approximation for rotation distance. *J. Graph Algorithms Appl.*, 14(2):385–390, 2010. `doi:10.7155/jgaa.00212`.

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

[20] S. Cleary and J. Taback. Bounding right-arm rotation distances. *Internat. J. Algebra Comput.*, 17(2):369–399, 2007. `doi:10.1142/S0218196707003664`.

[21] D. Comer. The ubiquitous B-tree. *ACM Comput. Surv.*, 11(2):121–137, 1979.

[22] K. Culik, II and D. Wood. A note on some tree similarity measures. *Inform. Process. Lett.*, 15(1):39–42, 1982. `doi:10.1016/0020-0190(82)90083-7`.

[23] M. Elder, E. Fusy, and A. Rechnitzer. Counting elements and geodesics in Thompson's group F. *J. Algebra*, 324(1):102–121, 2010. `doi:10.1016/j.jalgebra.2010.02.035`.

[24] D. B. A. Epstein, J. W. Cannon, D. F. Holt, S. V. F. Levy, M. S. Paterson, and W. P. Thurston. *Word processing in groups.* Jones and Bartlett Publishers, Boston, MA, 1992.

[25] S. B. Fordham. Minimal length elements of Thompson's group F. *Geom. Dedicata*, 99:179–220, 2003. `doi:10.1023/A:1024971818319`.

[26] S. B. Fordham and S. Cleary. Minimal length elements of Thompson's groups $F(p)$. *Geom. Dedicata*, 141:163–180, 2009. `doi:10.1007/s10711-008-9350-1`.

[27] D. Giampaolo. *Practical File System Design with the BE File System.* Morgan Kaufmann Publishers, 1999.

[28] D. Huguet. La structure du treillis des polyèdres de parenthésages. *Algebra Universalis*, 5:82–87, 1975. `doi:10.1007/BF02485235`.

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

[30] K. Leutola and J. Nieminen. Posets and generalized lattices. *Algebra Universalis*, 16(3):344–354, 1983. `doi:10.1007/BF01191789`.

[31] F. Luccio, A. Mesa Enriquez, and L. Pagli. k-restricted rotation distance between binary trees. *Inform. Process. Lett.*, 102(5):175–180, 2007. `doi:10.1016/j.ipl.2006.12.007`.

[32] J. Pallo. An efficient upper bound of the rotation distance of binary trees. *Inform. Process. Lett.*, 73(3-4):87–92, 2000. `doi:10.1016/S0020-0190(00)00008-9`.

[33] J. M. Pallo. The rotation χ-lattice of ternary trees. *Computing*, 66(3):297–308, 2001. `doi:10.1007/s006070170027`.

[34] J. M. Pallo. Right-arm rotation distance between binary trees. *Inform. Process. Lett.*, 87(4):173–177, 2003. `doi:10.1016/S0020-0190(03)00283-7`.

[35] L. Pournin. The diameter of associahedra. *Adv. Math.*, 259:13–42, 2014. `doi:10.1016/j.aim.2014.02.035`.

[36] O. Rodeh, J. Bacik, and C. Mason. Btrfs: The Linux B-tree filesystem. *ACM Trans. Storage*, 9(3):9:1–9:32, 2013. doi:10.1145/2501620.2501623.

[37] 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.

[38] D. D. Sleator, R. E. Tarjan, and W. P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1(3):647–681, 1988. doi:10.2307/1990951.

[39] M. Stein. Groups of piecewise linear homeomorphisms. *Trans. Amer. Math. Soc.*, 332(2):477–514, 1992. doi:10.2307/2154179.

[40] Z. Šunić. Tamari lattices, forests and Thompson monoids. *European J. Combin.*, 28(4):1216–1238, 2007. doi:10.1016/j.ejc.2006.02.001.

[41] C. Wladis. Cyclic subgroups are quasi-isometrically embedded in the Thompson-Stein groups. *Internat. J. Algebra Comput.*, 21(1-2):365–385, 2011. doi:10.1142/S0218196711006212.

[42] C. Wladis. Thompson's group is distorted in the Thompson-Stein groups. *Pacific J. Math.*, 250(2):473–485, 2011. doi:10.2140/pjm.2011.250.473.

[43] C. Wladis. The word problem and the metric for the Thompson-Stein groups. *J. Lond. Math. Soc. (2)*, 85(2):301–322, 2012. doi:10.1112/jlms/jdq060.