JOURNAL OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 1, Number 3, July 1988

ROTATION DISTANCE, TRIANGULATIONS, AND HYPERBOLIC GEOMETRY

DANIEL D. SLEATOR, ROBERT E. TARJAN, AND WILLIAM P. THURSTON

1. INTRODUCTION

A *rotation* in a binary tree is a local restructuring of the tree that changes it into another tree. One can execute a rotation by collapsing an internal edge of the tree to a point, thereby obtaining a node with three children, and then re-expanding the node of order three in the alternative way into two nodes of order 2. The *rotation distance* between a pair of trees is the minimum number of rotations needed to convert one tree into the other. The problem addressed in this paper is: what is the maximum rotation distance between any pair of n-node binary trees? We show that for all $n \geq 11$ this distance is at most $2n - 6$ and that for all sufficiently large n this bound is tight. Culik and Wood [2] showed that the maximum rotation distance is at most $2n - 2$. Tom Leighton (private communication) showed that there exist trees whose rotation distance is at least $7n/4 - O(1)$. Pallo [7] proposed a heuristic search algorithm to compute the rotation distance between two given trees.

Our interest in this problem stems from our attempt to solve the dynamic optimality conjecture concerning the performance of splaying [8, 10]. Splaying is a heuristic for modifying the structure of a binary search tree in such a way that repeatedly accessing and updating the information in the tree is efficient. Although our solution to the problem of maximum rotation distance did not resolve the conjecture about splaying, the results in this paper are interesting for at least two other reasons. First, the combinatorial system of trees and their rotations is a fundamental one that is isomorphic to other natural combinatorial

Received by the editors January 29, 1988.

1980 *Mathematics Subject Classification* (1985 *Revision*). Primary 51M20, 51M10, 05C05, 68P05.

This is a revised and expanded version of a paper that appeared in the 18th Annual ACM Symposium on Theory of Computing [9], Berkeley, CA, May 28–30, 1986.

Partial support provided for the first author by DARPA, ARPA order 4976, amendment 19, monitored by the Air Force Avionics Laboratory under contract F33615-87-C-1499, and by the National Science Foundation under grant CCR-8658139.

Partial support provided for the second author by the National Science Foundation under grant DCR-8605962.

Partial support provided for the third author by the National Science Foundation under grants DMR-8504984 and DCR-8505517.

systems. Results concerning this system are of interest from a purely mathematical point of view. Second, the method we use to solve the problem is novel and interesting in its own right and can potentially be applied to related problems.

A system that is isomorphic to binary trees related by rotations is that of triangulations of a polygon related by the *diagonal flip* operation. This is the operation that converts one triangulation of a polygon into another by removing a diagonal in the triangulation and adding the diagonal that subdivides the resulting quadrilateral in the opposite way. This type of move was studied by Wagner [14] in the context of arbitrary triangulated planar graphs and by Dewdney [3] in the case of graphs of genus one. They showed that any such graph can be transformed to any other by diagonal flips, but did not try to accurately estimate how many flips are necessary.

Our approach to solving the rotation distance problem is based on the observation that any sequence of diagonal flips converting one triangulation of a polygon into another gives a way to dissect (into tetrahedra) a polyhedron formed from the two triangulations. Using hyperbolic geometry, we construct polyhedra that require many tetrahedra to triangulate them. (Here and hereafter we use the word "triangulation" in a general sense meaning a dissection into simplices of appropriate dimension. A more rigorous definition is given in §2.4.) These polyhedra can be used to exhibit pairs of n-node trees (for all sufficiently large n) such that the rotation distance between them is $2n - 6$.

In §2 we define the problem on trees, make the connection between trees and triangulations of a polygon, and show that sequences of diagonal flips are related to triangulations of polyhedra. In §3 we show how to use hyperbolic geometry to obtain a lower bound on the number of tetrahedra required to triangulate any polyhedron. We then construct particular polyhedra that require many tetrahedra to triangulate them. §4 contains remarks and some open problems.

2. Definitions and equivalences

2.1. Binary trees. A binary tree is a collection of nodes of two types, external and internal, and three relations among these nodes: parent, left child, and right child. Every node except a special one called the root has a parent, and every internal node has a left and a right child. External nodes have no children. A tree is said to be of size n if it has n internal nodes. A tree of size n has $n+1$ external nodes. (See [5,10] for a more complete description of binary trees and tree terminology.) The number of steps required to walk from the root of the tree to a node is the depth of that node. (Each step moves from a node to one of its children.)

A *symmetric order* traversal of the tree visits all of the nodes exactly once. This order can be described by a recursive algorithm as follows: If the node is an internal node, traverse its left subtree in symmetric order, visit the node itself, then traverse its right subtree in symmetric order. If the node is an external node, merely visit it. The order in which the nodes are visited is called the

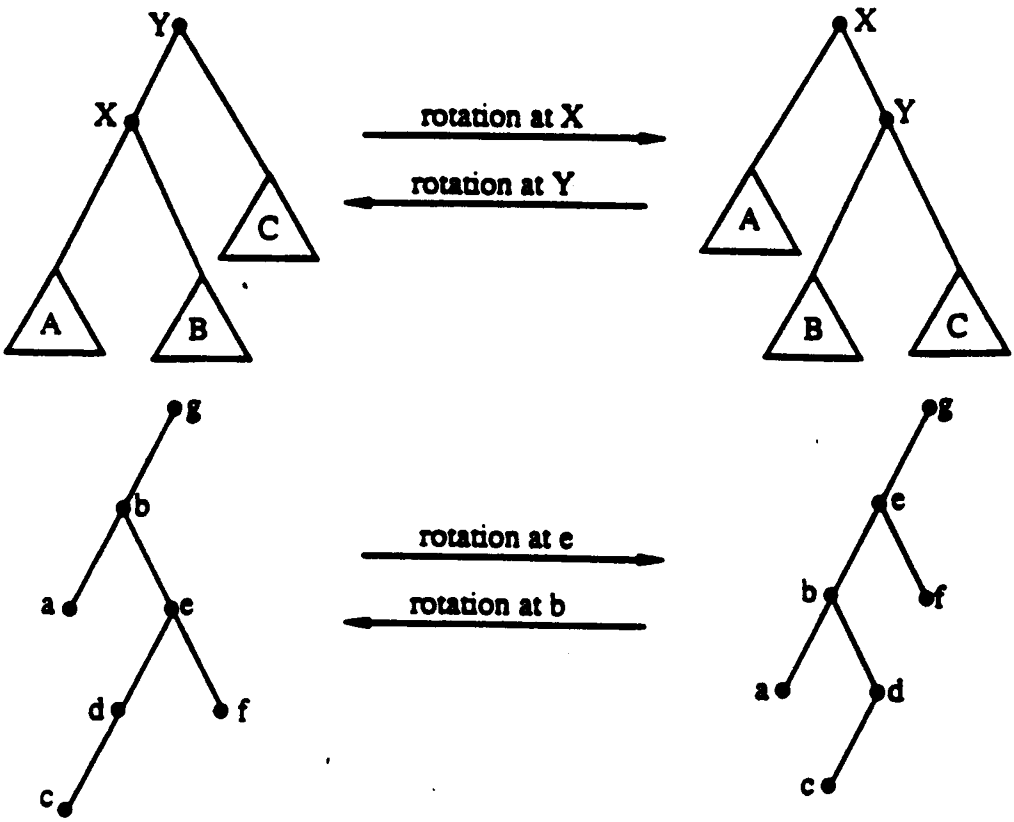

FIGURE 1. (a) The general definition of a rotation. Triangles denote subtrees. The tree shown could be part of a larger tree. (b) A rotation in a seven node tree. External nodes are not shown.

symmetric order permutation of the nodes (or simply the symmetric order of the nodes).

In a common computer-related application of binary trees the tree is used to store an ordered collection of pieces of information (called items). Each internal node of the tree is labeled with an item, and the order of the items is represented by the symmetric order of the nodes.

A *rotation* is an operation that changes one binary tree into another. In a tree of size n there are $n-1$ possible rotations, one corresponding to each nonroot internal node. Figure 1 shows the general rotation rule and the effect of a particular rotation on a particular tree. The rotation corresponding to a node changes the structure of the tree near that node, but leaves the structure elsewhere intact. A rotation maintains the symmetric order of the nodes, but changes the depths of some of them. Rotations are the primitives used by most schemes that maintain "balance" in binary trees [5, 10].

A rotation is an invertible operation; that is, if tree T can be changed into T' by a rotation, then T' can be changed back into T by a rotation. The *rotation graph* for trees of size n (denoted $RG(n)$) is an undirected graph with

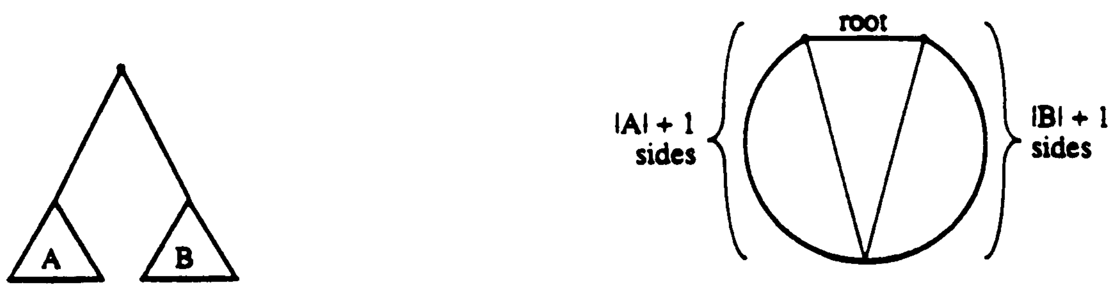

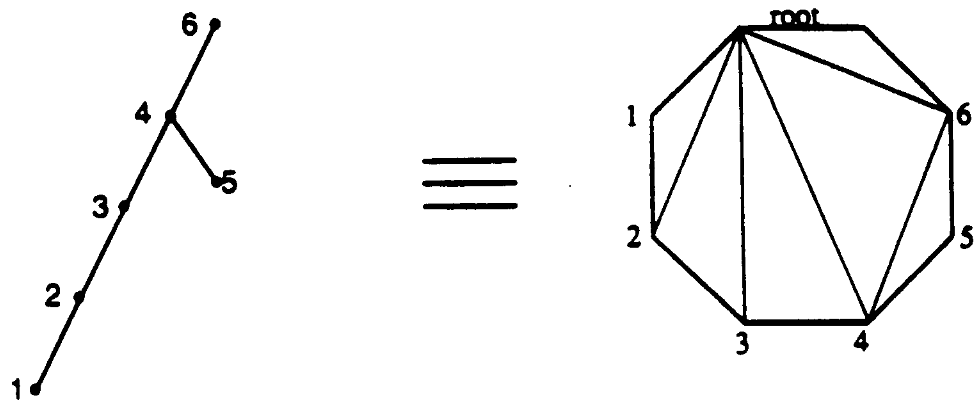

FIGURE 2. An example of a tree and its corresponding triangulation.

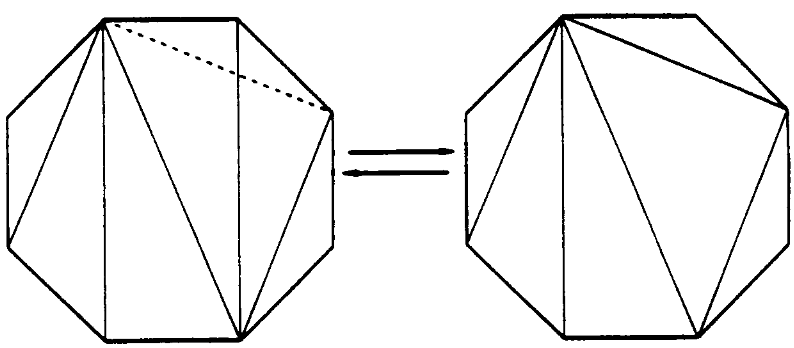

FIGURE 3. A diagonal flip in a triangulation of an octagon.

one vertex for each tree of size n and an edge between vertices T and T' if there is a rotation that changes T into T'.

Any binary tree of size n can be converted into any other by performing an appropriate sequence of rotations. Therefore the rotation graph is connected. We can define the *rotation distance* between two trees as the length of the shortest path in the rotation graph between the two trees, i.e. the minimum number of rotations required to convert one tree into the other. The main problem we address in this paper is that of estimating the diameter of $RG(n)$, i.e., the maximum rotation distance between any two n-node binary trees.