

An Efficient Algorithm for the Diameter of Cayley Graphs Generated by Transposition Trees

Ashwin Ganesan

Abstract—A problem of practical and theoretical interest is to determine or estimate the diameter of various families of Cayley networks. The previously known estimate for the diameter of Cayley graphs generated by transposition trees is an upper bound given in the oft-cited paper of Akers and Krishnamurthy (1989). In this work, we first assess the performance of their upper bound. We show that for every n , there exists a tree on n vertices, such that the difference between the upper bound and the true diameter value is at least $n - 4$.

Evaluating their upper bound takes time $\Omega(n!)$. In this paper, we provide an algorithm that obtains an estimate of the diameter, but which requires only time $O(n^2)$; furthermore, the value obtained by our algorithm is less than or equal to the previously known diameter upper bound. Such an improvement to polynomial time, while still performing at least as well as the previous bound, is possible because our algorithm works directly with the transposition tree on n vertices and does not require examining any of the permutations. We also provide a tree for which the value computed by our algorithm is not necessarily unique, which is an important result because such examples are quite rare. For all families of trees we have investigated so far, each of the possible values computed by our algorithm happens to also be an upper bound on the diameter.

Index Terms—Cayley graphs; transposition trees; diameter; algorithms; interconnection networks; permutations.

I. INTRODUCTION

A problem of practical and theoretical interest is to determine or estimate the diameter of various families of Cayley networks. In the field of interconnection networks, Cayley graphs generated by transposition trees were studied in Akers and Krishnamurthy [1], where it was shown that the diameter of some families of Cayley graphs is sublogarithmic in the number of vertices. This is one of the main reasons such Cayley graphs were considered a superior alternative to hypercubes for consideration as the topology of interconnection networks. Since then, much work has been done in this area; for further details, we refer the reader to [2], [3], [4], [5].

Let G be a group generated by a set of elements S . The Cayley graph (or Cayley diagram) of G with respect to the set of generators S , denoted by $\text{Cay}(G, S)$, is a directed

graph with vertex set G and with an arc from vertex g to vertex gs iff $g \in G$ and $s \in S$ (cf. [6], [7]). When S is closed under inverses, there is an arc from g to h if and only if there is an arc from h to g , and so $\text{Cay}(G, S)$ can be considered to be an undirected graph. When the identity group element I is not in S , the Cayley graph also has no self loops.

Cayley graphs possess a certain amount of symmetry. In a symmetric network, the topology of the network looks the same from every node. It is now known that many families of symmetric networks possess additional desirable properties such as optimal fault-tolerance [8], [9], algorithmic efficiency [10], optimal gossiping protocols [11] [12], and optimal routing algorithms [13], among others, and so have been widely studied in the fields of interconnection networks. New topologies continue to be proposed and assessed [14]. Cayley graphs, permutation groups, and distance-related questions continue to be an active area of research, and have also found applications in computational biology [15] and wireless sensor networks [16].

The diameter of a network represents the maximum communication delay between two nodes in the network. The design and performance of algorithms or bounds that determine or estimate the diameter of various families of Cayley graphs of permutation groups is of much theoretical and practical interest. The problem of determining the diameter of a Cayley network is the same as that of determining the diameter of the corresponding group for a given set of generators; the latter quantity is defined to be the minimum length of an expression for a group element in terms of the generators, maximized over all group elements. This diameter problem is difficult even for the simple case when the symmetric group is generated by cyclically adjacent transpositions [17]. The pancake flipping problem, which corresponds to determining the diameter of a particular permutation group, was studied in [18], and while some bounds were given there and an improvement was made recently [19], this problem remains open as well.

Throughout this paper, set of generators S will be a set of transpositions of $\{1, 2, \dots, n\}$. Given S , the transposition graph $T(S)$ is defined to be the simple, undirected graph whose vertex set is $\{1, 2, \dots, n\}$ and with vertices i and j being adjacent in $T(S)$ whenever $(i, j) \in S$. Cayley graphs of permutation groups generated by transposition sets include

Manuscript received June 11, 2012; revised October 4, 2012.

The author is with the Department of Mathematics, Amrita School of Engineering, Amrita University, Coimbatore 641112, Tamil Nadu, India.
Email: ashwin.ganesan@gmail.com

many topologies of interconnection networks that have been well studied; some special cases include the family of star graphs, bubble-sort graphs, modified bubble-sort graphs, and hypercubes, among others [3].

Algorithms for the diameter of Cayley graphs have been widely studied. When the generator set is fixed in advance, there do exist polynomial time algorithms for the problem of expressing a group element as a product of minimum length in terms of the generators and other distance-related problems (cf. [17], [20]). Since a Cayley graph is vertex-transitive, these results provide a polynomial time algorithm for the distance between any two vertices π and τ in the Cayley graph, since their distance is the minimum length generator sequence for $\pi^{-1}\tau$. However, these results only provide polynomial time algorithms for the distance between one given pair of vertices in the Cayley graph and not for the diameter of the entire Cayley graph, which is the maximum value of the distances between all pairs of vertices of the graph. Our focus in this paper is on the diameter problem.

A given set of transpositions S of $\{1, 2, \dots, n\}$ generates the entire symmetric group S_n iff the transposition graph $T(S)$ is connected [21]. A transposition graph which is a tree is called a transposition tree. Henceforth, S is a set of transpositions such that the transposition graph $T(S)$ is a tree. We often use the same symbol $T = T(S)$ to represent both the graph of the tree as well as a set of transpositions S , and the notation (i, j) is used to represent both an edge of T as well as the corresponding transposition in S . Since each transposition is its own inverse, the Cayley graph $\Gamma := \text{Cay}(S_n, S)$ is a simple, undirected graph. Let $\text{dist}_\Gamma(u, v)$ denote the distance between vertices u and v in an undirected graph Γ , and let $\text{diam}(\Gamma)$ denote the diameter of Γ . Note that $\text{dist}_\Gamma(\pi, \sigma) = \text{dist}_\Gamma(I, \pi^{-1}\sigma)$, where I denotes the identity permutation. Thus, the diameter of Γ is the maximum of $\text{dist}_\Gamma(I, \pi)$ over $\pi \in S_n$.

In [1] (cf. also [22, p. 188]), it is shown that the diameter of $\Gamma := \text{Cay}(S_n, S)$ is bounded as

$$\text{diam}(\Gamma) \leq \max_{\pi \in S_n} \left\{ c(\pi) - n + \sum_{i=1}^n \text{dist}_T(i, \pi(i)) \right\},$$

where the maximum is over all permutations in S_n , $c(\pi)$ denotes the number of cycles in the disjoint cycle representation of π , and dist_T is the distance function on pairs of vertices of the tree.

Observe that evaluating this upper bound requires $\Omega(n!)$ computations since the quantity in braces above needs to be evaluated for each permutation in S_n . When a bound or algorithm is proposed in the literature, it is often of interest to determine how far away the bound can be from the true value in the worst case, and to obtain more efficient algorithms for estimating the parameters. In this paper, we assess the performance of the previously known upper bound on the

diameter, propose a new $O(n^2)$ algorithm to estimate the diameter of Cayley graphs for any given transposition tree, and we investigate the properties and performance of our algorithm.

II. PRELIMINARIES, AND SUMMARY OF OUR MAIN RESULTS

We now recall the previous relevant results and some terminology from the literature as well as summarize our contributions to this problem.

Let S_n denote the set of all permutations of $\{1, 2, \dots, n\}$. We represent a permutation $\pi \in S_n$ as a linear arrangement, as in $[\pi(1), \pi(2), \dots, \pi(n)]$, or in cycle notation. $c(\pi)$ denotes the number of cycles in π , including cycles of length 1. Thus, if $\pi = [3, 5, 1, 4, 2] = (1, 3)(2, 5) \in S_5$, then $c(\pi) = 3$. For $\pi, \tau \in S_n$, $\pi\tau$ is the permutation obtained by applying τ first and then π . If $\pi \in S_n$ and $\tau = (i, j)$ is a transposition, then $c(\tau\pi) = c(\pi) + 1$ if i and j are in the same cycle of π , and $c(\tau\pi) = c(\pi) - 1$ if i and j are in different cycles of π (cf. [23]). $\text{Fix}(\pi)$ denotes the set of fixed points of π . We assume throughout that the transposition tree has at least 5 vertices since the problem is easily solved for all smaller trees by using brute force.

Throughout this work, Γ denotes the Cayley graph $\text{Cay}(S_n, S)$ generated by a transposition tree $T = T(S)$. The previous bound on the diameter is as follows:

Theorem 1. [1] *Let $\Gamma := \text{Cay}(S_n, S)$ be the Cayley graph generated by a transposition tree $T(S)$. Then, for any $\pi \in S_n$,*

$$\text{dist}_\Gamma(I, \pi) \leq c(\pi) - n + \sum_{i=1}^n \text{dist}_T(i, \pi(i)),$$

where $c(\pi)$ is the number of cycles (including fixed points) in the disjoint cycle representation of π .

Since Γ is vertex-transitive and $\text{dist}_\Gamma(\pi, \tau) = \text{dist}_\Gamma(I, \pi^{-1}\tau)$, by taking the maximum over both sides of the above inequality, we obtain:

Corollary 2. [22, p.188]

$$\text{diam}(\Gamma) \leq \max_{\pi \in S_n} \left\{ c(\pi) - n + \sum_{i=1}^n \text{dist}_T(i, \pi(i)) \right\} =: f(T).$$

In the sequel we shall often refer to $f(T)$ as the the previously known upper bound on the diameter of the Cayley graph or *the diameter upper bound*. Note that evaluating this estimate $f(T)$ requires time $\Omega(n!)$ since the quantity in braces needs to be evaluated for each of the $n!$ vertices of the Cayley graph. While one can investigate methods to optimize such an algorithm, at present there is no known polynomial time algorithm for computing $f(T)$. It is not known whether an algorithm for computing $f(T)$ can be optimized to run

in polynomial time or time better than $\Omega(n!)$. Our objective is to take a different approach and propose a new algorithm that works directly with the transposition tree on n vertices rather than the Cayley graph on $n!$ vertices.

We now recall from [1] a proof sketch of these results since we refer to this terminology in the sequel. Suppose we are given a transposition tree T on vertex set $\{1, 2, \dots, n\}$ and a permutation $\pi \in S_n$ for which we wish to determine $\text{dist}_\Gamma(I, \pi)$. At each vertex i of the tree, we place a marker labeled $\pi(i)$. Thus, the permutation π represents the current position of the markers $1, 2, \dots, n$ on the tree. To *apply an edge* (i, j) of the tree to the current position of markers is to say that we switch the markers at the endpoints of the edge (i, j) . Note that the permutation corresponding to the new position of the markers is exactly $\pi(i, j)$ (here, we read products or compositions of permutations from right to left). The problem of determining $\text{dist}_\Gamma(I, \pi)$ is thus equivalent to that of determining the minimum number of edges necessary to ‘home’ each marker i to vertex i of the tree, and the Cayley graph Γ represents the state transition diagram of the current position of markers. This problem of placing markers $\pi(i)$ at vertex i of an n -vertex graph and homing each marker to its vertex is sometimes also referred to as ‘sorting’ a permutation using only the transpositions defined by T . Such problems also arise in routing problems in Cayley networks; for example, a node σ receiving a message destined to node τ of Γ , or equivalently, a node $\pi = \tau^{-1}\sigma$ receiving a message destined to node I , needs to figure out which of its neighbors in Γ is closest to the destination node, and this amounts to determining which edge of the tree is optimal in terms of the objective of sorting the current permutation of markers using the minimum number of edges. Further details can be found in [1].

Let T be a tree on vertex set $\{1, \dots, n\}$, and let $\pi \neq I$ be the current position of markers on the tree. It can be shown that T always has an edge ij such that the edge satisfies one of the following two conditions: Either (A) the marker at i and the marker at j will both reduce their distance to $\pi(i)$ and $\pi(j)$, respectively, if the edge (i, j) is applied, or (B) the marker at one of i or j is already homed, and the other marker wishes to apply the edge (i, j) . Let $f_T(\pi)$ denote the upper bound quantity in the right hand side of the inequality in Theorem 1. Thus, $f(T) = \max_{\pi \in S_n} f_T(\pi)$. It can be shown that during each step that a transposition corresponding to an edge of type A or type B is applied to π , we get a new position of markers π' which has a strictly smaller value of f_T ; i.e., $f_T(\pi') < f_T(\pi)$, and it can be verified that $f_T(I) = 0$. This proves the bounds above.

We point out that this same diameter upper bound inequality of Corollary 2 is also derived in [24]; however, this paper was published in 1991, whereas Akers and Krishnamurthy [1] was published in 1989 and widely picked up on in

the interconnection networks community by then. There are some subsequent papers, such as [25] and [26], which cite only [24] and not [1].

Note that the distance and diameter bounds above need not hold if T has cycles (the proof recalled above breaks down because if T has cycles, there exists a $\pi \neq I$ such that T has no admissible edges for this π). Thus, when we study the strictness of the diameter upper bound, we assume that T is a tree and Γ is the Cayley graph generated by this tree.

The exact diameter value of Cayley graphs generated by transposition trees is known in only some special cases. For example, if the transposition tree is a path graph on n vertices, the corresponding Cayley graph is called a bubble-sort graph. It is well known that the diameter of this Cayley graph is equal to the maximum number of inversions of a permutation, which is $\binom{n}{2}$ (cf. [1] [27]). When the transposition tree is a star $K_{1, n-1}$, the Cayley graph is called a star graph, and it has diameter equal to $\lceil 3(n-1)/2 \rceil$ (cf. [1]). For the general case of arbitrary trees, only bounds such as Corollary 2 are known.

It is possible to obtain a heuristic derivation of the diameter upper bound formula $f(T)$, as follows. It is straightforward to derive the distance upper bound for the special case when the transposition tree is a star $K_{1, n-1}$, and we get [1]

$$\text{dist}_\Gamma(I, \pi) \leq n + c(\pi) - 2|\text{Fix}(\pi)| - r(\pi).$$

Observe that $|\text{Fix}(\pi)| = n - |\overline{\text{Fix}(\pi)}|$, which yields

$$\text{dist}_\Gamma(I, \pi) \leq c(\pi) - n + 2|\overline{\text{Fix}(\pi)}| - r(\pi).$$

Note that when the tree is a star, $2|\overline{\text{Fix}(\pi)}|$ is almost (i.e. within 1 of) the sum of distances $\sum_{i=1}^n \text{dist}_T(i, \pi(i))$. This leads us to the question of whether the inequality

$$\text{dist}_\Gamma(I, \pi) \leq c(\pi) - n + \sum_{i=1}^n \text{dist}_T(i, \pi(i))$$

also holds for all the remaining trees T , and this question has been answered affirmatively by Theorem 1.

We shall later use the following result on the sharpness of the diameter upper bound inequality:

Theorem 3. [28] *Let Γ denote the Cayley graph generated by a transposition tree T . Then the diameter upper bound inequality*

$$\text{diam}(\Gamma) \leq f(T)$$

holds with equality if T is a path.

The main results of this paper are as follows.

It is of interest to know how far away the diameter upper bound $f(T)$ bound can be from the true value in the worst case. We show that for every n , there exists a transposition tree on n vertices such that the difference between the

diameter upper bound and the true diameter value of the Cayley graph is at least $n - 4$. This result gives a lower bound on the difference, and we leave it as an open problem to determine an upper bound for this difference.

We provide a new algorithm (Algorithm A below) which more efficiently computes, for any given transposition tree, an estimate of the diameter of the Cayley graph generated by the tree. Remarkably, the proposed algorithm requires only time $O(n^2)$ to compute, whereas no polynomial time algorithm is known for computing the previous bound $f(T)$. Furthermore, it is shown that the value computed by Algorithm A is at least as good as (i.e. is less than or equal to) the previous upper bound $f(T)$. Such a result is possible because the new algorithm works directly with the transposition tree on n vertices and does not require examining the vertices of the Cayley graph; it is only the proofs of our results that require examining the individual permutations. It is proved that sometimes the value obtained by Algorithm A is strictly better than (i.e. is strictly less than) the previous upper bound in the literature (cf. Theorem 8). Some advantages of the new algorithm over the previous upper bound are illustrated; for example, the proofs related to the worst case performance (of $n - 4$) of the new algorithm are much simpler than those of the previous upper bound (cf. Proposition 9 and the remarks preceding it). It is also shown that the value computed Algorithm A is not necessarily unique (Theorem 7); this is an important result because such counterexamples are quite rare.

For all families of trees we have investigated so far, each of the possible values β computed by Algorithm A is an upper bound on the diameter, i.e.

$$\text{diam}(\Gamma) \leq \beta \leq f(T);$$

here, we prove that the second inequality holds for all trees, and the first inequality holds for many families of trees (in fact for all trees investigated so far).

Some further interesting questions and open problems on this algorithm and related bounds are discussed towards the end of this paper.

III. STRICTNESS OF THE DIAMETER UPPER BOUND

Recall that the diameter of a Cayley graph Γ generated by a transposition tree T is bounded as

$$\text{diam}(\Gamma) \leq \max_{\pi \in S_n} \left\{ c(\pi) - n + \sum_{i=1}^n \text{dist}_T(i, \pi(i)) \right\} =: f(T).$$

We now assess the performance of this bound and derive a strictness result. The results in this section also appear in the conference paper [29].

Define the worst case performance of this upper bound by the quantity

$$\Delta_n := \max_{T \in \mathcal{T}_n} |f(T) - \text{diam}(\Gamma)|,$$




Fig. 1. A transposition tree T on n vertices.

where \mathcal{T}_n denotes the set of all trees on n vertices.

Theorem 4. For every $n \geq 5$, there exists a tree $T = T(S)$ on n vertices such that the difference between the actual diameter of the Cayley graph $\text{diam}(\text{Cay}(S_n, S))$ and the diameter upper bound $f(T)$ is at least $n - 4$; in other words, $\Delta_n \geq n - 4$.

Proof: Throughout this proof, we let T denote the transposition tree defined by the edge set $\{(1, 2), (2, 3), \dots, (n - 3, n - 2), (n - 2, n - 1), (n - 2, n)\}$, which is shown in Figure 1. For conciseness, we let $d(i, j)$ denote the distance in T between vertices i and j . Also, for leaf vertices i, j of T , we let $T - \{i, j\}$ denote the tree on $n - 2$ vertices obtained by removing vertices i and j of T .

Our proof is in two parts. In the first part we establish that $f(T)$ is equal to $\binom{n}{2} - 2$. In the second part we show that the diameter of the Cayley graph generated by T is at most $\binom{n-1}{2} + 1$. Together, this yields the desired result.

We now present the first part of the proof; we establish that $f(T)$, defined by

$$f(T) := \max_{\sigma \in S_n} \left\{ c(\sigma) - n + \sum_{i=1}^n \text{dist}_T(i, \sigma(i)) \right\},$$

is equal to $\binom{n}{2} - 2$. We prove this result by examining several sub-cases. Define

$$f_T(\sigma) := c(\sigma) - n + S_T(\sigma), \quad S_T(\sigma) := \sum_{i=1}^n \text{dist}_T(i, \sigma(i)).$$

We consider two cases, (1) and (2), depending on whether 1 and n are in the same or different cycle of σ ; each of these cases will further involve subcases. In most of these subcases, we show that for a given σ , there is a σ' such that $f_T(\sigma) \leq f_T(\sigma')$ and $f_T(\sigma') \leq \binom{n}{2} - 2$.

(1) Assume 1 and n are in the same cycle of σ . So $\sigma = (1, k_1, \dots, k_s, n, j_1, \dots, j_\ell)\hat{\sigma}$. The different subcases consider the different possible values for s and ℓ .

(1.1) Suppose $s = 0, \ell = 0$. So $\sigma = (1, n)\hat{\sigma} = (1, n)\sigma_2 \dots \sigma_r$. Then, $f_T(\sigma) = c(\sigma) - n + S_T(\sigma) = r - n + 2(n - 2) + S_{T - \{1, n\}}(\hat{\sigma}) = 2n - 5 + (r - 2) + (n - 2) + S_{T - \{1, n\}}(\hat{\sigma}) = 2n - 5 + c(\hat{\sigma}) + (n - 2) + S_{T - \{1, n\}}(\hat{\sigma}) = 2n - 5 + f_{T - \{1, n\}}(\hat{\sigma}) \leq 2n - 5 + \binom{n-2}{2} = \binom{n}{2} - 2$, where by Theorem 3 the inequality holds with equality for some




Fig. 2. Positions of j_1 and j_ℓ arising in subcase (1.4.1).

$\hat{\sigma}$. Thus, the maximum of $f_T(\sigma)$ over all permutations that contain $(1, n)$ as a cycle is equal to $\binom{n}{2} - 2$. It remains to show that for all other kinds of permutations σ in the symmetric group S_n , $f_T(\sigma) \leq \binom{n}{2} - 2$.

(1.2) Suppose $s = 1, \ell = 0$. So $\sigma = (1, i, n)\sigma_2 \dots \sigma_r = (1, i, n)\hat{\sigma}$. We consider some subcases.

(1.2.1) Suppose $i = n - 1$. Then, $f_T(\sigma) = r - n + (2n - 2) + S_{T-\{1, n-1, n\}}(\hat{\sigma}) = 2n - 4 + f_{T-\{1, n-1, n\}}(\hat{\sigma}) \leq 2n - 4 + \binom{n-3}{2} \leq \binom{n}{2} - 2$, where the inequality is by Theorem 3.

(1.2.2) Suppose $2 \leq i \leq n - 2$; so $\sigma = (1, i, n)\hat{\sigma}$. Let $\sigma' = (1, n)(i)\hat{\sigma}$. It is easily verified that $f_T(\sigma) \leq f_T(\sigma')$, and so the desired bound follows from applying subcase (1.1) to $f_T(\sigma')$.

(1.3) Suppose $s = 0, \ell = 1$, so $\sigma = (1, n, i)\hat{\sigma}$. Since $f_T(\sigma) = f_T(\sigma^{-1})$, this case also is settled by (1.2).

(1.4) Suppose $s = 0, \ell \geq 2$, so $\sigma = (1, n, j_1, \dots, j_\ell)\hat{\sigma}$. Let $\sigma' = (1, n)(j_1, \dots, j_\ell)\hat{\sigma}$. Observe that $f_T(\sigma) \leq f_T(\sigma')$ iff $d(n, j_1) + d(j_\ell, 1) \leq d(n, 1) + d(j_\ell, j_1) + 1$. We prove the latter inequality by considering 4 subcases:

(1.4.1) Suppose $j_1 < j_\ell \leq n - 2$. Then, an inspection of the tree in Figure 2 shows that $d(n, j_1) + d(j_\ell, 1) = d(n, 1) + d(j_\ell, j_1)$, and so the inequality holds.

(1.4.2) Suppose $j_1 > j_\ell$ and $j_1, j_\ell \leq n - 2$. Then, $d(n, j_1) + d(j_\ell, 1) \leq d(1, n)$, and so the inequality holds.

(1.4.3) Suppose $j_1 = n - 1$. Then $d(n, j_1) = 2$. Also, $d(j_\ell, 1) \leq d(n, 1)$ and $d(j_\ell, j_1) \geq 1$, and so the inequality holds.

(1.4.4) Suppose $j_\ell = n - 1$. Then, $d(j_\ell, 1) = d(n, 1)$ and $d(n, j_1) = d(j_\ell, j_1)$, and so again the inequality holds.

(1.5) Suppose $s = 1, \ell = 1$, so $\sigma = (1, i, n, j)\hat{\sigma}$. Let $\sigma' = (1, n)(i, j)\hat{\sigma}$.

(1.5.1) If $i = n - 1$, by symmetry in T between vertices n and $n - 1$, this subcase is resolved by subcase (1.4).

(1.5.2) Let $2 \leq i \leq n - 2$. Then $d(1, i) + d(i, n) = d(1, n)$. So $f_T(\sigma) \leq f_T(\sigma')$ iff $d(n, j) + d(j, 1) \leq d(1, n) + d(i, j) + d(j, i) + 1$, which is true since $d(n, j) + d(j, 1) \leq d(1, n) + 2$.

(1.6) Suppose $s = 1, \ell \geq 2$. So $\sigma = (1, i, n, j_1, \dots, j_\ell)\hat{\sigma}$. Let $\sigma' = (1, n)(i, j_1, \dots, j_\ell)\hat{\sigma}$. It suffices to show that $f_T(\sigma) \leq f_T(\sigma')$, i.e., that $d(1, i) + d(i, n) + d(n, j_1) + d(j_\ell, 1) \leq d(1, n) + d(1, n) + d(i, j_1) + d(j_\ell, i) + 1$. We examine the terms of this latter inequality for various subcases:

(1.6.1) Suppose $2 \leq i \leq n - 2$. Then $d(1, i) + d(i, n) = d(1, n)$.

(1.6.1a) If $j_\ell = n - 1$, then $d(j_\ell, i) = n - i - 1$ and $d(i, j_1) = |i - j_1|$, and so the inequality holds iff $-1 \leq |j_1 - i| + j_1 - i$, which is clearly true.

(1.6.1b) Suppose $2 \leq j_\ell \leq n - 2$. Then, the inequality holds iff $d(n, j_1) + j_\ell - 1 \leq n - 2 + |i - j_1| + |i - j_\ell| + 1$, which can be verified separately for the cases $j_1 = n - 1$ and $2 \leq j_1 \leq n - 2$.

(1.6.2) Suppose $i = n - 1$. By symmetry in T of the vertices n and $n - 1$, this case is resolved by (1.4).

(1.7) Suppose $s \geq 2, \ell = 0$, so $\sigma = (1, k_1, \dots, k_s, n)\hat{\sigma}$. Since $f_T(\sigma) = f_T(\sigma^{-1})$, this case is resolved by (1.4).

(1.8) Suppose $s \geq 2, \ell = 1$, so $\sigma = (1, k_1, \dots, k_s, n, j_1)\hat{\sigma}$. Let $\sigma' = (1, n)(k_1, \dots, k_s, j_1)\hat{\sigma}$. We can assume that $2 \leq k_1 \leq n - 2$ since the $k_1 = n - 1$ case is resolved by (1.4) due to the symmetry in T . To show $f_T(\sigma) \leq f_T(\sigma')$, it suffices to prove the inequality $d(1, k_1) + d(k_s, n) + d(n, j_1) + d(j_1, 1) \leq d(1, n) + d(n, 1) + d(k_s, j_1) + d(j_1, k_1) + 1$. We prove this inequality by separately considering whether $j_1 = n - 1$ or $k_s = n - 1$ or neither:

(1.8.1) Suppose $j_1 = n - 1$. Substituting $d(k_s, n) = n - k_s - 1, d(n, j_1) = 2, d(j_1, 1) = j_1 - 1$, etc, we get that the inequality holds iff $k_1 \leq |j_1 - k_1| + n - 2$, which is clearly true.

(1.8.2) Suppose $2 \leq j_1 \leq n - 2$. Then $d(n, j_1) = n - j_1 - 1$, and so the inequality holds iff $k_1 + d(k_s, n) + n - 3 \leq 2n - 3 + d(k_s, j_1) + d(j_1, k_1)$. If $k_s = n - 1$, this reduces to $j_1 + k_1 \leq 2n - 3 + |j_1 - k_1|$, and is true, whereas if $2 \leq k_s \leq n - 2$, this reduces to $k_1 - k_s \leq 1 + |j_1 - k_s| + |j_1 - k_1|$, which is true due to the triangle inequality.

(1.9) Suppose $s, l \geq 2$, so $\sigma = (1, k_1, \dots, k_s, n, j_1, \dots, j_\ell)\hat{\sigma}$.

Let $\sigma' = (1, k_1, \dots, k_s, n)(j_1, \dots, j_\ell)\hat{\sigma}$. It suffices to show that $S_T(\sigma) \leq S_T(\sigma') + 1$, i.e., that $d(n, j_1) + j_\ell \leq n + d(j_1, j_\ell)$.

(1.9.1) If $j_1 < j_\ell$, then $j_1 \leq n - 2$, and so $d(n, j_1) = n - j_1 - 1$ and $d(j_1, j_\ell) = j_\ell - j_1$; the inequality thus holds.

(1.9.2) If $j_1 > j_\ell$, then $d(j_1, j_\ell) = j_1 - j_\ell$, and so it suffices to show that $d(n, j_1) \leq n + j_1 - 2j_\ell$. It can be verified that this holds if $j_1 = n - 1$ and also if $2 \leq j \leq n - 2$.

(2) Now suppose 1 and n are in different cycles of σ . So let $\sigma = (1, k_1, \dots, k_s)(n, j_1, \dots, j_\ell)\hat{\sigma}$.

(2.1) Suppose $s = 0$. Then $f_T(\sigma) \leq \binom{n-1}{2} - 2$, by induction on n .

(2.2) Suppose $s = 1$. So let $\sigma = (1, i)(n, j_1, \dots, j_\ell)\hat{\sigma}$. By symmetry in T between vertices n and $n - 1$ and subcase (1.1), we may assume $i \neq n - 1$. Let $\sigma' = (1, n)(i, j_1, \dots, j_\ell)\hat{\sigma}$. It suffices to show that $S_T(\sigma) \leq S_T(\sigma')$. If $\ell = 0$ this is clear since $d(1, i) \leq d(1, n)$. Suppose $\ell \geq 2$. Then, by the triangle inequality, $d(n, j_1) + d(j_\ell, n) \leq d(j_1, i) + d(i, n) + d(i, j_\ell) + d(i, n) = d(j_1, i) + d(i, j_\ell) + (n - i - 1)2$. Also, $d(1, n) = d(1, i) + d(i, n) = d(1, i) + n - i - 1$.

Hence, $2d(1, i) + d(n, j_1) + d(j_\ell, n) \leq 2d(1, n) + d(i, j_1) + d(j_\ell, i)$. Hence, $S_T(\sigma) \leq S_T(\sigma')$. The case $\ell = 1$ can be similarly resolved by substituting j_1 for j_ℓ in the $\ell \geq 2$ case here.

(2.3) Suppose $s \geq 2, \ell = 0$. Then, by Theorem 3, $f_T(\sigma) \leq \binom{n-1}{2}$.

(2.4) Suppose $s \geq 2, \ell = 1$, so $\sigma = (1, k_1, \dots, k_s)(n, j_1)\hat{\sigma}$. Let $\sigma' = (1, n)(k_1, \dots, k_s, j_1)\hat{\sigma}$. It suffices to show that $d(1, k_1) + d(1, k_s) + 2d(n, j_1) \leq 2d(1, n) + d(k_s, j_1) + d(k_1, j_1)$. This inequality is established by considering the two subcases:

(2.4.1) Suppose $j_1 = n - 1$. Then the inequality holds iff $2k_s + k_1 \leq 3n - 7 + |k_1 - j_1|$, which is true since $k_1, k_2 \leq n - 2$ and $|k_1 - j_1| \geq 1$.

(2.4.2) Suppose $j_1 \neq n - 1$. Then the inequality holds iff $k_1 - j_1 + k_s - j_1 \leq |k_1 - j_1| + |k_s - j_1|$, which is clearly true.

(2.5) Suppose $s, \ell \geq 2$, so $\sigma = (1, k_1, \dots, k_s)(n, j_1, \dots, j_\ell)\hat{\sigma}$.

Let $\sigma' = (1, n)(k_1, \dots, k_s, j_1, \dots, j_\ell)\hat{\sigma}$. To show $f_T(\sigma) \leq f_T(\sigma')$, it suffices to show that $d(1, k_1) + d(k_s, 1) + d(n, j_1) + d(j_\ell, n) \leq 2d(n, 1) + d(k_s, j_1) + d(j_\ell, k_1)$. By symmetry in T between vertices n and $n - 1$, we may assume $k_1, \dots, k_s \neq n - 1$, since these cases were covered in (1). We establish this inequality as follows:

(2.5.1) Suppose $j_1 = n - 1$. Then $d(n, j_1) = 2$ and $d(n, j_\ell) = n - j_\ell - 1$. So the inequality holds iff $2k_s \leq 2(n - 2) + |j_\ell - k_1| + j_\ell - k_1$, which is true since $k_s \leq n - 2$ and $|j_\ell - k_1| + j_\ell - k_1 \geq 0$.

(2.5.2) Suppose $j_1 \neq n - 1$. Then $d(n, j_1) = n - j_1 - 1$. If $j_\ell = n - 1$, the inequality holds iff $2k_1 \leq 2(n - 2) + j_1 - k_s + |j_1 - k_s|$, which is true since $k_1 \leq n - 2$. If $j_\ell \neq n - 1$, the inequality holds iff $k_s - j_1 + k_1 - j_\ell \leq |k_s - j_1| + |k_1 - j_\ell|$, which is true.

This concludes the first part of the proof.

We now provide the second part of the proof. Let Γ be the Cayley graph generated by T . We show that $\text{diam}(\Gamma) \leq \binom{n-1}{2} + 1$. Let $\pi \in S_n$, and suppose each vertex i of T has marker $\pi(i)$. We show that all markers can be homed using at most the proposed number of transpositions. Since $\text{diam}(T) = n - 2$, marker 1 can be moved to vertex 1 using at most $n - 2$ transpositions. Now remove vertex 1 from the tree T , and repeat this procedure for marker 2, and then for marker 3, and so on, removing each vertex from T after its marker is homed. Continuing in this manner, we eventually arrive at a star $K_{1,3}$, whose Cayley graph has diameter 4. Hence, the diameter of Γ is at most $[(n - 2) + (n - 3) + \dots + 5 + 4 + 3] + 4 = \binom{n-1}{2} + 1$. This completes the proof. ■

Let $s(n)$ denote the number of non-isomorphic trees on n vertices. Let Δ_n be the strictness as defined above. Then, computer simulations yield the results in Table I:

TABLE I
STRICTNESS OF THE DIAMETER UPPER BOUND

n	5	6	7	8	9
$s(n)$	3	6	11	23	47
Δ_n	1	2	3	4	6

These results imply that the $n - 4$ lower bound for Δ_n is best possible, and an open problem is to obtain an upper bound for Δ_n . That this lower bound is exact for $n = 5$ can also be obtained using the results given above, as follows. There are only three nonisomorphic trees on 5 vertices, namely the tree given in the proof of Theorem 4, for which the $n - 4$ lower bound is achieved, and the trees of maximum diameter (the path) and minimum diameter (the star), for which the diameter bound $f(T)$ is known to be exact (cf. [28]).

IV. THE ALGORITHM

We now provide an algorithm that takes as its input a transposition tree $T(S)$ on n vertices and provides as output an estimate of the diameter of the Cayley graph $\text{Cay}(S_n, S)$. The notation used to describe our algorithm should be self-explanatory and is similar to that used in Knuth [30].

Algorithm A

Given a transposition tree $T = T(S)$, this algorithm computes a value β which is an estimate for the diameter of the Cayley graph $\text{Cay}(S_n, S)$. $|V(T)|$ denotes the current value of the number of vertices in T ; initially, $V(T) = \{1, 2, \dots, n\}$.

A1. [Initialize.]

Set $\beta \leftarrow 0$.

A2. [Find two vertices i, j of T that are a maximum distance apart.]

Find any two vertices i, j of T such that $\text{dist}_T(i, j) = \text{diam}(T)$.

A3. [Update β , and remove i, j from T .]

Set $\beta \leftarrow \beta + (2 \text{diam}(T) - 1)$, and set $T \leftarrow T - \{i, j\}$. If T still has 3 or more vertices, return to step A2; otherwise, set $\beta \leftarrow \beta + |V(T)| - 1$ and terminate this algorithm. ■

Example 1. Consider the transposition tree $\{(1, 2), (2, 3), (3, 4), (4, 5), (4, 6), (3, 7), (7, 8)\}$ shown in Figure 3. If Algorithm A picks the sequence of vertex pairs during step A2 to be $\{1, 8\}, \{5, 7\}$ and $\{2, 6\}$, then the value returned by the algorithm is $\beta = 7 + 5 + 5 + 1 = 18$. On the other hand, if Algorithm A picks the vertex pairs to be $\{1, 5\}, \{6, 8\}$ and $\{2, 7\}$, then the value returned by the algorithm is still $\beta = 7 + 7 + 3 + 1 = 18$. In this example, the value returned by the algorithm is unique even though the subtrees $T - \{1, 8\}$ and $T - \{1, 5\}$ are non-isomorphic and even have different diameters. ■




Fig. 3. The transposition tree on 8 vertices in Example 1.

Despite the outcome in the above example where the value computed by Algorithm A is independent of the vertex pairs chosen in step A2, it is shown later that there do exist trees for which the value computed by the algorithm is not unique (i.e. the final value depends on which vertex pairs were chosen during step A2), though this non-uniqueness property is rare.

The eccentricity of a vertex u in a graph is defined to be the maximum value of the distance from u to a vertex of the graph. The center of a graph is defined to be the set of vertices of minimum eccentricity. It is known that the center of a tree is either a single vertex or two adjacent vertices. Also, every path of maximum length in a tree passes through its center.

In view of this, one way to implement step A2, which picks any two vertices of the tree that are a maximum distance apart, is as follows. Start with an arbitrary vertex u of the tree, and do a depth-first search to find a vertex i farthest that is from u (i and u will be on different ‘sides’ of the center). Then start at vertex i and do another depth-first search to find a vertex j that is farthest from i . Then, the resulting $i - j$ path has maximum length in the tree. (Alternatively, one could also carry out a breadth-first search rather than a depth-first search.) Since a depth-first search on the transposition tree on n vertices takes time $O(n)$, the first iteration of Algorithm A takes time $O(n)$, the second iteration takes time $O(n - 2)$, and so on. Thus, Algorithm A takes time $O(n^2)$. The previously known estimate in the literature for the diameter of Cayley graphs generated by transposition trees works on each of the $n!$ vertices of the Cayley graph and takes time $\Omega(n!)$ to evaluate; hence, Algorithm A, is a significant improvement over the previous bound.

V. PROPERTIES AND PERFORMANCE OF THE ALGORITHM

In this section we prove that the value obtained by Algorithm A is less than or equal to the previously known diameter upper bound. We also show that there exist instances where the diameter estimate computed by Algorithm A is strictly better than (i.e. is strictly less than) than the

previously known diameter upper bound $f(T)$. Furthermore, we construct a rare transposition tree for which the value computed by Algorithm A is not unique.

Theorem 5. *Let T be a transposition tree on vertex set $\{1, 2, \dots, n\}$, and let β be the value obtained by Algorithm A for this tree. Then, β is less than or equal to the previously known upper bound on the diameter of the Cayley graph, i.e.*

$$\beta \leq f(T).$$

Proof: Let $\{i_1, j_1\}, \{i_2, j_2\}, \dots, \{i_r, j_r\}$ be the vertex pairs chosen by Algorithm A during the r iterations of step A2, where $r = \lfloor (n - 1)/2 \rfloor$. Construct a permutation π as follows. If n is odd, then T contains only one vertex, i_{r+1} say, when the algorithm terminates. In this case, we let $\pi = (i_1, j_1) \dots (i_r, j_r)(i_{r+1}) \in S_n$. If n is even, then T contains two vertices, i_{r+1} and j_{r+1} say, when the algorithm terminates. In this case, we let $\pi = (i_1, j_1) \dots (i_{r+1}, j_{r+1}) \in S_n$. In either case, $r + 1 = \lceil n/2 \rceil$, the value of β computed by the algorithm equals

$$\beta = \left(\sum_{\ell=1}^r \{2 \text{dist}_T(i_\ell, j_\ell) - 1\} \right) + \{(n + 1) \pmod 2\},$$

and the quantity $f_T(\pi) := c(\pi) - n + \sum_{i=1}^n \text{dist}_T(i, \pi(i))$ evaluates to

$$f_T(\pi) = (r+1)n - n + \left(2 \sum_{\ell=1}^r \text{dist}_T(i_\ell, j_\ell) \right) + 2 \{(n + 1) \pmod 2\}.$$

A quick check shows that the expressions above for β and $f_T(\pi)$ are equal. Hence, for every sequence of vertex pairs chosen by Algorithm A, there exists a permutation π such that the value β returned by Algorithm A is at most $f_T(\pi)$. Hence, $\beta \leq \max_{\pi \in S_n} f_T(\pi)$. ■

Since each of the possible values β computed by Algorithm A is at most the previous upper bound $f(T)$, it follows immediately that the largest of the possible values computed by Algorithm A, denoted by β_{\max} , is also at most the previous upper bound. We now show that β_{\max} is an upper bound on the diameter the Cayley graph:

Theorem 6. *Let Γ be the Cayley graph generated by a transposition tree T . Let β_{\max} denote the maximum possible value returned by Algorithm A for this tree. Then,*

$$\text{diam}(\Gamma) \leq \beta_{\max} \leq f(T).$$

Proof: The second inequality has already been proved. We now prove the first inequality. Let $\pi \in S_n$. Suppose that initially each vertex k of the tree has marker $\pi(k)$. It suffices to show that all markers can be homed to their respective vertices using at most β_{\max} edges of the tree.

Consider the following procedure. Pick any two vertices i, j of T that are a maximum distance apart. We consider

two cases, depending on the distance in T between vertex i and the current location $\pi^{-1}(i)$ of the marker i :

Case 1: Suppose that the distance in T between vertices i and $\pi^{-1}(i)$ is at most $\text{diam}(T) - 1$. Then marker i can be homed using at most $\text{diam}(T) - 1$ transpositions. And then, marker j can be homed using at most $\text{diam}(T)$ edges. Hence, markers i and j can both be homed to leaf vertices i and j , respectively, using at most $2 \text{diam}(T) - 1$ edges. We now let $i_1 = i$ and $j_1 = j$.

Case 2: Now consider the case where the distance in T between vertices i and $\pi^{-1}(i)$ is equal to $\text{diam}(T)$. Let x be the unique vertex of the tree adjacent to $\pi^{-1}(i)$. In the first sequence of steps, marker $\pi^{-1}(i)$ can be homed to vertex $\pi^{-1}(i)$ using at most $\text{diam}(T)$ edges. The last of these transpositions will home marker $\pi^{-1}(i)$ and place marker i at vertex x , whose distance to i is exactly $\text{diam}(T) - 1$. In the second sequence of steps, marker i can be homed to vertex i using at most $\text{diam}(T) - 1$ edges. Hence, using these two sequences of steps, markers i and $\pi^{-1}(i)$ can both be homed using at most $2 \text{diam}(T) - 1$ edges. We now let $i_1 = i$ and $j_1 = \pi^{-1}(i)$.

We now remove from T the vertices i_1 and j_1 , and repeat this procedure on $T - \{i_1, j_1\}$ to get another pair $\{i_2, j_2\}$. Continuing in this manner until T contains at most two vertices, we see that all markers can be homed using at most

$$\left\{ \sum_{\ell=1}^r (2 \text{dist}_T(i_\ell, j_\ell) - 1) \right\} + \{(n + 1) \bmod 2\}$$

edges. This quantity is equal to the value β returned by the Algorithm when it chooses $\{i_1, j_1\}, \dots, \{i_r, j_r\}$ as its vertex pairs during each iteration of step A2, and hence this quantity is at most β_{\max} . Thus, $\text{dist}_\Gamma(I, \pi) \leq \beta_{\max}$ for all $\pi \in S_n$. ■

Let \mathcal{B} denote the set of possible values that can be the output of Algorithm A. Thus, $\beta_{\max} := \max_{\beta \in \mathcal{B}} \beta$. An open problem is to determine whether *each* of the possible values returned by the algorithm is an upper bound on the diameter, i.e. whether β is an upper bound on the diameter of the Cayley graph for each $\beta \in \mathcal{B}$. The examples studied so far show that for many families of trees (in fact, for all the trees investigated so far), the minimum possible value returned by the algorithm is also an upper bound on the diameter. Thus, we believe that each of the possible values $\beta \in \mathcal{B}$ is an upper bound on the diameter of the Cayley graph, but we do not have a proof of this.

We now prove that the value returned by the algorithm is not necessarily unique:

Theorem 7. *There exist transposition trees for which the value returned by Algorithm A is not unique.*

Proof: Consider the transposition tree defined by the edge set $\{(1, 2), (2, 3), (3, 6), (3, 4), (4, 5), (6, 7)$,




Fig. 4. The transposition tree used in the proof of Theorem 7 and Theorem 8.

$(6, 8), (6, 9)\}$, which is shown in Figure 4. If Algorithm A picks the sequence of vertex pairs during step A2 to be $\{1, 5\}, \{2, 7\}, \{4, 8\}$ and $\{3, 9\}$, then the value returned by the algorithm is $\beta = 7 + 5 + 5 + 3 = 20$. And if Algorithm A picks the vertex pairs to be $\{1, 7\}, \{5, 8\}, \{2, 9\}$ and $\{4, 6\}$, then the value returned by the algorithm is $\beta = 7 + 7 + 5 + 3 = 22$. Hence, \mathcal{B} contains $\{20, 22\}$. ■

In addition to time complexity, another advantage of Algorithm A over the previous upper bound is that the value computed by Algorithm A can sometimes be *strictly* less than the previous upper bound on the diameter.

Theorem 8. *The value computed by Algorithm A is always less than or equal to the previously known upper bound on the diameter, and there exist transposition trees for which the value computed by Algorithm A is strictly less than the previous upper bound.*

Proof: The first part of the assertion has been proved earlier. For the second part, consider the transposition tree shown in Figure 4. It can be confirmed with the help of a computer that the true diameter value of the Cayley graph generated by this tree T is 18 and the previous upper bound $f(T)$ on the diameter evaluates to 22. As mentioned above, if Algorithm A picks the sequence of vertex pairs during step A2 to be $\{1, 5\}, \{2, 7\}, \{4, 8\}$ and $\{3, 9\}$, then the value returned by the algorithm is $\beta = 7 + 5 + 5 + 3 = 20$, which is strictly less than the previous upper bound. ■

In Theorem 4 (cf. also the recent work Ganesan [29]), it was shown that the difference between the previous upper bound $f(T)$ and the actual diameter $\text{diam}(\Gamma)$ is at least $n - 4$. The proof given there is quite involved and required an examination of several (over 25) subcases.

As another advantage of Algorithm A over the previous upper bound $f(T)$, we show that the value computed by Algorithm A can also have a difference of at least $n - 4$ from the actual diameter value but that the proof of this result is much simpler than the corresponding result for the previous upper bound $f(T)$:

Proposition 9. *For every n , there exists a transposition tree*




Fig. 5. The transposition tree used in the proof of Proposition 9

on n vertices, such that the difference between the value computed by Algorithm A and the actual diameter value of the Cayley graph is at least $n - 4$.

Proof: Consider the transposition tree $\{(1, 2), (2, 3), \dots, (n - 3, n - 2), (n - 2, n - 1), (n - 2, n)\}$ shown in Figure 5. The diameter of this tree is $n - 2$. After Algorithm A picks and removes two vertices from this tree that are a distance $n - 2$ apart, we obtain the path graph on $n - 2$ vertices. The unique value computed by Algorithm A is thus $\{2(n - 2) - 1\} + \{2(n - 3) - 1\} + \{2(n - 5) - 1\} + \dots$, which equals $\{2(n - 2) - 1\} + \binom{n-2}{2} = \binom{n-1}{2} + n - 3$.

Now, any permutation can be sorted on this tree using at most $\binom{n-1}{2} + 1$ edges. Indeed, marker 1 can be homed to its vertex using at most $n - 2$ edges, and this vertex can then be removed from the tree. Marker 2 can then be homed using at most $n - 3$ edges, and so on, and marker $n - 4$ can be homed using at most 3 edges. At this point, we arrive at a star $K_{1,3}$, and any permutation on this star can be sorted using at most 4 edges since the diameter of the Cayley graph generated by this star is equal to 4. Thus, the diameter of the Cayley graph generated by this tree is at most $(n - 2) + (n - 3) + \dots + 3 + 4 = \binom{n-1}{2} + 1$.

Hence, for this transposition tree, the difference between the value computed by Algorithm A and the actual diameter value is at least $n - 4$. ■

Note that Theorem 6 implies that the value computed by Algorithm A is an upper bound on the diameter for all trees for which $|\mathcal{B}| = 1$ since for such trees any value β computed by the algorithm is equal to β_{\max} . For such trees, Algorithm A efficiently computes a value which is both an upper bound on the diameter as well as better than (or at least as good as) the previously known diameter upper bound. In the rare case when the value computed by Algorithm A is not unique, it still seems likely that each of the possible values computed by the algorithm is an upper bound on the diameter, though we do not have a proof for this.

VI. CONCLUDING REMARKS

Cayley graphs have been studied as a suitable model for the topology of interconnection networks, and a problem of both theoretical and practical interest is to obtain bounds for the diameter of Cayley graphs. In this work, we investigated

an upper bound on the diameter of Cayley graphs generated by transposition trees. We showed that for every $n > 4$, there exists a tree on n vertices such that the difference between the previous upper bound $f(T)$ and the true diameter is at least $n - 4$. Such results are of interest because they give us insight as to how far away these bounds can be from the true diameter value in the worst case and sometimes tell us for which families of graphs this bound can be utilized or not utilized.

The $n - 4$ lower bound on Δ_n was seen to be best possible in the sense that it is attained for some values of n . Now consider the tree on 9 vertices consisting of the edge set $\{(1, 2), (2, 3), (3, 4), (4, 5), (5, 6), (6, 7), (6, 8), (6, 9)\}$. Then, it can be confirmed with the help of a computer that the diameter of the Cayley graph generated by this tree is 24, and the diameter upper bound $f(T)$ for this tree evaluates to 30. Hence, this $n - 4$ lower bound is not the exact value of the strictness Δ_n , and an open problem is to obtain an upper bound for Δ_n .

An efficient algorithm (Algorithm A) for the diameter of Cayley graphs generated by transposition trees was given. Remarkably, the algorithm has time complexity $O(n^2)$, compared to the previous upper bound in the literature $f(T)$ for which no polynomial time algorithm is known and for which the existing methods take time $\Omega(n!)$ to compute. We proved that the value obtained by our algorithm is less than or equal to the previously known diameter upper bound. Sometimes the value computed by the algorithm is *strictly* less than the previous upper bound. Such an improvement in efficiency from time $\Omega(n!)$ to polynomial time $O(n^2)$, while still performing at least as well as the previous bound, was possible because we worked directly with the transposition tree on n vertices, and so our algorithm does not require examining any of the permutations.

We described some further advantages of our algorithm over the previous bound, besides the improvement in time complexity. We provided a tree for which the value computed by the algorithm is not necessarily unique. This is an important fact because such counterexamples are quite rare.

We believe that each of the possible values computed by Algorithm A is an upper bound on the diameter of the Cayley graph, but we do not have a proof for this. For the families of trees investigated so far, the maximum possible value returned by our algorithm is exactly equal to the previously known diameter upper bound. However, our algorithm arrived at the same value using a very different (and also simpler and more efficient) method than the previously known diameter upper bound, and investigating further properties of this algorithm might lead to new insights on this problem.

The algorithm presented here raises many further interesting questions and problems. For example, is it true that the

algorithm returns a unique value (i.e. $|\mathcal{B}| = 1$) for almost all trees? Is it true that for all trees, the maximum possible value β_{\max} returned by the algorithm is equal to $f(T)$? Other open problems include to characterize those trees for which the value returned by Algorithm A is unique, and to characterize those trees for which the sequences of subtrees generated by the algorithm are isomorphic (i.e. are independent of the choice of vertex pairs during step A2, unlike the tree in Example 1 above).

REFERENCES

- [1] S. B. Akers and B. Krishnamurthy, "A group-theoretic model for symmetric interconnection networks," *IEEE Transactions on Computers*, vol. 38, no. 4, pp. 555–566, 1989.
- [2] S. Lakshminarayanan, J.-S. Jho, and S. K. Dhall, "Symmetry in interconnection networks based on Cayley graphs of permutation groups: A survey," *Parallel Computing*, vol. 19, pp. 361–407, 1993.
- [3] M. C. Heydemann, "Cayley graphs and interconnection networks," in *Graph symmetry: algebraic methods and applications (Editors: Hahn and Sabidussi)*. Kluwer Academic Publishers, Dordrecht, 1997, pp. 167–226.
- [4] D. Z. Du and F. Hsu, *Combinatorial Network Theory*. Applied Optimization Series, Springer, 1995.
- [5] J. Xu, *Topological Structure and Analysis of Interconnection Networks*. Kluwer Academic Publishers, 2001.
- [6] B. Bollobás, *Modern Graph Theory*. Graduate Texts in Mathematics vol. 184, Springer, New York, 1998.
- [7] N. L. Biggs, *Algebraic Graph Theory, 2nd Edition*. Cambridge University Press, Cambridge, 1993.
- [8] B. Alspach, "Cayley graphs with optimal fault tolerance," *IEEE Transactions on Computers*, vol. 41, no. 10, pp. 1337–1339, 1992.
- [9] S. B. Akers and B. Krishnamurthy, "On group graphs and their fault tolerance," *IEEE Transactions on Computers*, vol. 36, no. 7, pp. 885–888, 1987.
- [10] F. Annexstein, M. Baumslag, and A. L. Rosenberg, "Group action graphs and parallel architectures," *SIAM Journal on Computing*, vol. 19, no. 3, pp. 544–569, 1990.
- [11] J. C. Bermond, T. Kodate, and S. Perennes, "Gossiping in Cayley graphs by packets," in *Proceedings of the Franco-Japanese conference Brest July 95*, vol. Lecture Notes in Computer Science, 1120. Springer verlag, 1996, pp. 301–315.
- [12] S. Zhou, "A class of arc-transitive Cayley graphs as models for interconnection networks," *SIAM Journal on Discrete Mathematics*, vol. 23, no. 2, pp. 694–714, 2009.
- [13] B. Chen, W. Xiao, and B. Parhami, "Internode distance and optimal routing in a class of alternating group networks," *IEEE Transactions on Computers*, vol. 55, no. 12, pp. 1645–1648, 2006.
- [14] B. Parhami, "Swapped interconnection networks: Topological, performance, and robustness attributes," *Journal of Parallel and Distributed Computing*, vol. 65, no. 11, pp. 1443–1452, 2005.
- [15] V. Bafna and P. A. Pevzner, "Sorting by transpositions," *SIAM Journal on Computing*, vol. 11, no. 2, pp. 224–240, 1998.
- [16] L. Wang and K. W. Tang, "The Cayley Graph implementation in TinyOS for dense wireless sensor networks," in *Proc. of the 6th Wireless Telecommunications Symposium (WTS'07)*, 2007.
- [17] M. Jerrum, "The complexity of finding minimum length generator sequences," *Theoretical Computer Science*, vol. 36, pp. 265–289, 1985.
- [18] W. H. Gates and C. H. Papadimitriou, "Bounds for sorting by prefix reversal," *Discrete Mathematics*, vol. 27, pp. 47–57, 1979.
- [19] B. Chitturi, W. Fahle, Z. Meng, L. Morales, C. O. Shields, I. H. Sudborough, and W. Voit, "An $(18/11)n$ upper bound for sorting by prefix reversals," *Theoretical Computer Science*, vol. 410, no. 36, pp. 3372–3390, 2009.
- [20] G. Cooperman and L. Finkelstein, "New methods for using cayley graphs in interconnection networks," *Discrete Applied Mathematics*, vol. 37, pp. 95–118, 1992.
- [21] C. Godsil and G. Royle, *Algebraic Graph Theory*. Graduate Texts in Mathematics vol. 207, Springer, New York, 2001.
- [22] G. Hahn and G. Sabidussi, *Graph Symmetry: Algebraic Methods and Applications*. Kluwer Academic Publishers, Dordrecht, 1997.
- [23] N. L. Biggs, *Discrete Mathematics, 2nd Edition*. Oxford University Press, 2003.
- [24] T. P. Vaughan, "Bounds for the rank of a permutation on a tree," *Journal of Combinatorial Mathematics and Combinatorial Computing*, vol. 19, pp. 65–81, 1991.
- [25] T. P. Vaughan and F. J. Portier, "An algorithm for the factorization of permutations on a tree," *Journal of Combinatorial Mathematics and Combinatorial Computing*, vol. 18, pp. 11–31, 1995.
- [26] J. H. Smith, "Factoring, into edge transpositions of a tree, permutations fixing a terminal vertex," *Journal of Combinatorial Theory Series A*, vol. 85, pp. 92–95, 1999.
- [27] C. Berge, *Principles of Combinatorics*. Academic Press, New York, 1971.
- [28] A. Ganesan, "On the sharpness of a bound on the diameter of Cayley graphs generated by transposition trees," *Journal of Combinatorial Mathematics and Combinatorial Computing*, to appear.
- [29] —, "On the strictness of a bound for the diameter of Cayley graphs generated by transposition trees," in *Proceedings of the International Conference on Mathematical Modeling and Scientific Computation*, ser. Communications in Computer and Information Science, vol. 283. Gandhigram, India. Springer-Verlag, 2012, pp. 54–61.
- [30] D. Knuth, *Combinatorial Algorithms*, ser. The Art of Computer Programming. Addison-Wesley Professional, 2011, vol. 4A.

Ashwin Ganesan was born in Kalpakkam, Tamil Nadu, India, in 1977. He received the B.S. degree in electrical engineering from Marquette University, Milwaukee, WI, USA, in 1998, and the M.S. degree in electrical engineering from the University of Wisconsin at Madison, WI, USA, in 2000. His research interests are in combinatorics, graph theory, algorithms, discrete mathematics, and their applications.

He received the Top Scholar in Curriculum Award from Marquette University in 1998, the Frank Rogers Bacon Fellowship from the University of Wisconsin at Madison, WI, USA during 1998-1999, and the Regents' Fellowship from the University of California at Berkeley, CA, USA, during 2001-2002. He was a Teaching and Research Assistant at the University of California at Berkeley during 2001-2004 and at the University of Wisconsin at Madison during 2004-2008. He was a Senior Lecturer and Assistant Professor at Mumbai University affiliated engineering colleges during 2008-2010. Since 2010, he has been an Assistant Professor in the Department of Mathematics, Amrita School of Engineering, Amrita University, Tamil Nadu, India.