On the sharpness of a bound on the diameter of Cayley graphs generated by transposition trees

ASHWIN GANESAN
Department of Mathematics
Amrita School of Engineering
Amrita Vishwa Vidyapeetham
Amritanagar, Coimbatore-641112, India.
e-mail: ashwin.ganesan@gmail.com

Abstract

Let Γ be a Cayley graph generated by a transposition tree T on n vertices. In an oft-cited paper [1] (see also [9]), it was shown that the diameter of the Cayley graph Γ on $n!$ vertices is bounded as

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

where the maximization is over all permutations π in S_n, $c(\pi)$ denotes the number of cycles in π, and dist_T is the distance function in T. It is of interest to determine for which families of trees this inequality holds with equality. In this work, we first investigate the sharpness of this upper bound. We prove that the above inequality is sharp for all trees of maximum diameter (i.e. all paths) and for all trees of minimum diameter (i.e. all stars), but the bound can still be strict for trees that are non-extremal. We also show that a previously known inequality on the distance between vertices in some families of Cayley graphs holds with equality and we prove that for some families of graphs an algorithm related to these bounds is optimal.

Keywords: Cayley graphs, permutation groups, diameter, transposition trees, permutations, interconnection networks.
2010 Mathematics Subject Classification Number: 05C85, 05C25.

1 Introduction

Let Γ be a Cayley graph on $n!$ vertices generated by a transposition tree on the vertex set $\{1, 2, \ldots, n\}$. In an oft-cited paper [1], an upper bound was provided for the diameter of Γ, in terms of distances in the underlying transposition tree. In that paper, the authors also showed that the diameter of the Cayley graph generated by a star $K_{1,n-1}$ is $\lfloor 3(n-1)/2 \rfloor$, which is sublogarithmic in the number of vertices of the Cayley graph, whereas the diameter of the n-dimensional hypercube on 2^n vertices is n, which is logarithmic in the number of vertices. Thus, Cayley graphs generated by transposition trees were since considered an attractive alternative to hypercubes as a consideration for the topology of interconnection networks. Cayley graphs also offer other desirable properties such as optimal fault tolerance, optimal gossiping protocols [4], and optimal algorithmic efficiency [2], among other properties.

Let S denote a set of transpositions of $\{1, 2, \ldots, n\}$. We can describe S by its transposition graph $T(S)$, which is a simple, undirected graph whose vertex set is $\{1, 2, \ldots, n\}$ and with vertices i and j being adjacent whenever $(i, j) \in S$. Construct the Cayley graph $\Gamma = Cay(Gr(S), S)$, whose vertex set is the permutation group generated by S and with two vertices g and h being adjacent in Γ if and only if there exists an $s \in S$ such that $gs = h$ [5],[7].

A transposition graph which is a tree is called a transposition tree. When $T(S)$ contains a spanning tree, S generates the entire symmetric group on n letters [3]. Let $T = T(S)$ denote the transposition tree corresponding to the set of transpositions S. Since each element of S is its own inverse, we assume Γ is a simple, undirected graph. Let $\mathrm{dist}_G(u, v)$ denote the distance between vertices u and v in an undirected graph G, and let $\mathrm{diam}(G)$ denote the diameter of G. Note that $\mathrm{dist}_\Gamma(\pi, \sigma) = \mathrm{dist}_\Gamma(I, \pi^{-1}\sigma)$. Thus, the diameter of Γ is the maximum of $\mathrm{dist}_\Gamma(I, \pi)$ over $\pi \in S_n$.

A natural problem is to understand how the properties of the Cayley graph Γ depend on the properties of the underlying transposition graph $T(S)$. Our question here concerns specifically the diameter of the Cayley graph as well as the distances between vertices of the Cayley graph, which we wish to express in terms of parameters of the underlying transposition graph. There are some upper bounds in the literature for this specific problem [1], [10, p.188]. In this work we investigate these bounds further.

It is of interest to determine for what families of graphs the previously known inequalities for the diameter and distances between vertices of the Cayley graph are sharp, and the extent to which these bounds can be away from the true diameter value in the worst case. In this work, we investigate the sharpness of these bounds. We show that the bound is sharp (i.e. the inequality holds with equality) for all families of trees of maximum diameter (i.e. for all paths) and for all trees of minimum diameter (i.e. for all stars). On the other side, we exhibit some trees for which this bound is strict. We also investigate the optimality of an algorithm posed in [1] for estimating the distance between vertices in the Cayley graph, and we prove that for some families of trees this algorithm sorts any given permutation using the minimum number of transpositions.

1.1 Notation and terminology

Let S_n denote the symmetric group on $[n] := \{1, 2, \ldots, n\}$. We represent a permutation $\pi \in S_n$ as an arrangement of $[n]$, as in $[\pi(1), \pi(2), \ldots, \pi(n)]$ or in cycle notation. $c(\pi)$ denotes the number of cycles in π, including cycles of length 1. Also, $\mathrm{inv}(\pi)$ denotes the number of inversions of π (cf. [3]). Thus, if $\pi = [3, 5, 1, 4, 2] = (1, 3)(2, 5) \in S_5$, then $c(\pi) = 3$ and $\mathrm{inv}(\pi) = 6$. 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 part of the same cycle of π, and $c(\tau\pi) = c(\pi) - 1$ if i and j are in different cycles of π; and similarly for $c(\pi\tau)$ (cf. [6]). I denotes the identity element of S_n. $\mathrm{Fix}(\pi)$ denotes the set of fixed points of π, and $\overline{\mathrm{Fix}(\pi)}$ denotes the complement set $[n] - \mathrm{Fix}(\pi)$.

Throughout this work, Γ denotes the Cayley graph generated by a transposition tree T. We now recall some previously known bounds. We also outline the proof of these bounds since we refer to the proof in the sequel.

Theorem 1.1. [1] *Let T be a tree and let $\pi \in S_n$. Let Γ be the Cayley graph generated by T. Then*

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

By taking the maximum over both sides, it follows that

Corollary 1.2. [10, p.188]

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

In the sequel, we refer to the first upper bound as the *distance upper bound* $f_T(\pi)$ and the second upper bound as the *diameter upper bound* $f(T)$:

Definition 1.3. For a *transposition tree* T and $\pi \in S_n$, define

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

We now recall from [1] the proofs of these results since we refer to the proof in the sequel. Start with a given tree T on vertices labeled by $[n]$ and an element $\pi \in S_n$, $\pi \neq I$, for which we wish to determine $\text{dist}_\Gamma(I, \pi)$. Initially, at each vertex i of T, we place a marker $\pi(i)$. Multiplying π by the transposition (i, j) amounts to *switching* the markers at vertices i and j. We now have a new set of markers at each vertex of T corresponding to the permutation $\pi(i, j)$, which is a vertex adjacent to π in Γ. The problem of determining $\text{dist}_\Gamma(I, \pi)$ is then equivalent to that of finding the minimum number of switches necessary to home each marker (i.e. to bring each marker i to vertex i). Given any T with vertex set $[n]$ and markers for these vertices corresponding to $\pi \neq I$, 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 switch (i, j) is applied, or (B) the marker at one of i or j is already homed, and the other marker wishes to use the switch (i, j). We call an edge that satisfies one of these two conditions an *admissible edge* of type A or type B. It can be shown that during each step that a transposition τ corresponding to an admissible edge is applied to π, we get a new vertex π' which has a strictly smaller value of the left hand side above; i.e., $f_T(\pi') < f_T(\pi)$, and it can be verified that $f_T(I) = 0$. This proves the bounds above. This algorithm, which we call the AK algorithm (after Akers and Krishnamurthy), can be viewed as 'sorting' a permutation using only the transpositions defined by T, and the Cayley graph is the state transition diagram of the current permutation of markers.

The diameter of Cayley graphs generated by transposition trees is known for some particular families of graphs. For example, if the transposition tree is a path graph on n vertices, then the diameter of the corresponding Cayley graph is the maximum number of inversions of a permutation, which is $\binom{n}{2}$, and if the transposition tree is a star $K_{1,n-1}$, then the diameter of the corresponding Cayley graph is $\lfloor 3(n-1)/2 \rfloor$ (cf. [1]). For the special case when T is a star, another upper bound on the distance between vertices in the Cayley graph is known:

Lemma 1.4. [1] *Let T be a star. Then*

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

where $r(\pi)$ equals 0 if $\pi(1) = 1$ and $r(\pi) = 2$ otherwise (here, the center vertex of T is assumed to have the label 1).

Note that the distance and diameter bounds above need not hold if T has cycles (the proof mentioned 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 sharpness (or lack thereof) of the upper bounds, we assume throughout that T is a tree and Γ is the Cayley graph generated by a tree.

We point out that this same diameter upper bound inequality is also derived in Vaughan [12]; 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 [13] and [11], which cite only Vaughan [12] and not [1].

1.2 Summary of our contributions

In this work, we investigate the performance of the upper bounds given in [1] for the distances between vertices of a Cayley graph and for the diameter of Cayley graphs generated by transposition trees. We now summarize our contributions here:

Let Γ denote the Cayley graph generated by a transposition tree T. We show that the previously known distance upper bound

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

is exact or sharp (i.e. the inequality holds with equality) for all $\pi \in S_n$ if and only if T is a star.

We also show that the previously known diameter upper bound

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

is sharp if T is a star or a path. Note that this also implies that even though the distance upper bound is *not* exact for any paths, if we take the maximum over both sides, we get that the diameter upper bound *is* exact for all paths (i.e. the strict inequality becomes an equality when we maximize over all $\pi \in S_n$).

It was shown in [1] that: when T is a star,

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

We show here that this inequality holds with equality.

Enroute to deriving the diameter upper bound, the authors in [1] proposed an algorithm, which we called the AK algorithm above, for finding the distance between vertices in the Cayley graph. We prove some properties about this algorithm in Section 3.

We conclude with some open problems and further directions.

2 Sharpness of the distance and diameter upper bounds

In our proofs, it will be convenient to define

$$S_T(\pi) := \sum_{i=1}^{n} \mathrm{dist}_T(i, \pi(i)).$$

Thus, $f_T(\pi) = c(\pi) - n + S_T(\pi)$. While the bounds and results here are independent of the labeling of the vertices of T, it will be convenient to assume that the center vertex of a star has label 1, and that the vertices of a path are labeled consecutively from 1 to n. Also, note that the diameter and distance bounds are invariant to a translation of the labels on the set of integers, i.e. we can replace the labels $\{1, 2, \dots, n\}$ by say $\{2, 3, \dots, n+1\}$.

Theorem 2.1. *Let Γ be the Cayley graph generated by a transposition tree T. Then, the distance upper bound inequality*

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

holds with equality for all $\pi \in S_n$ if and only if T is the star $K_{1,n-1}$.

Proof. Suppose T is the star $K_{1,n-1}$. It is already known that $\mathrm{dist}_\Gamma(I, \pi) \leq f_T(\pi)$ for all $\pi \in S_n$. We now prove the reverse inequality. We want to show that $f_T(\pi)$ is the minimum number of transpositions of the form $(1, i), 2 \leq i \leq n$ required to sort π. Each vertex i of T is initially assigned the marker $\pi(i)$. Before the markers along edge $(1, i)$ are interchanged, there are four possibilities for the values of the marker $\pi(1)$ at vertex 1 and marker $\pi(i)$ at vertex i:

(a) $\pi(1) = 1$ and $\pi(i) = i$: In this case, applying transposition $(1, i)$ creates a new permutation π' for which $c(\pi)$ has reduced by 1 (i.e. $c(\pi') = c(\pi) - 1$) and $S_T(\pi)$ has increased by 2, thereby increasing $f_T(\pi)$ by 1.

(b) $\pi(1) = 1$ and $\pi(i) \neq i$: Applying transposition $(1, i)$ reduces $c(\pi)$ and doesn't affect $S_T(\pi)$. Hence, $f_T(\pi)$ is reduced by 1.

(c) $\pi(1) \neq 1$ and $\pi(i) = i$: Applying $(1, i)$ reduces $c(\pi)$ by 1 and increases $S_T(\pi)$ by 2, hence increases $f_T(\pi)$ by 1.

(d) $\pi(1) \neq 1$ and $\pi(i) \neq i$: There are four subcases here:

(d.1) The first subcase is when $\pi(1) = i$ and $\pi(i) = 1$. In this case, applying $(1, i)$ increases $c(\pi)$ by 1, and reduces $S_T(\pi)$ by 2, thereby reducing $f_T(\pi)$ by 1.

(d.2) Suppose $\pi(1) = j$ (where $j \neq 1, i$) and $\pi(i) = 1$. Applying $(1, i)$ increases $c(\pi)$ by 1 and doesn't affect $S_T(\pi)$. So $f_T(\pi)$ increases by 1.

(d.3) Suppose $\pi(1) = i$ and $\pi(i) = j \neq 1$. Then applying $(1, i)$ increases $c(\pi)$ by 1 and reduces $S_T(\pi)$ by 2, and hence reduces $f_T(\pi)$ by 1.

(d.4) Suppose $\pi(1) = k$ and $\pi(i) = j$, where $j, k \neq 1, i$. Then applying $(1, i)$ changes $c(\pi)$ by 1 and doesn't change $S_T(\pi)$, so that $f_T(\pi)$ changes by 1.

In all cases above, switching the markers on an edge $(1, i)$ of T reduces $f_T(\pi)$ by at most 1. Hence, the minimum number of transpositions required to sort π, or equivalently, the value of $\text{dist}_\Gamma(I, \pi)$, is at least $f_T(\pi)$. Hence, $\text{dist}_\Gamma(I, \pi) \geq f_T(\pi)$ for all $\pi \in S_n$. This proves the reverse inequality.

Observe that $S_T(\pi) = 2|\overline{\text{Fix}(\pi)}|$ when $\pi(1) = 1$, and $S_T(\pi) = 2|\overline{\text{Fix}(\pi)}| - 2$ otherwise. Thus, it is seen that when T is the star graph, $f_T(\pi)$ evaluates to $c(\pi) - n + 2|\overline{\text{Fix}(\pi)}| - r(\pi)$. Hence, $\text{dist}_\Gamma(I, \pi) = c(\pi) - n + 2|\overline{\text{Fix}(\pi)}| - r(\pi)$.

Now suppose T is not a star. Then, $\text{diam}(T) \geq 3$. So T contains 4 ordered vertices i, j, k and ℓ that comprise a path of length 3. Let π be the permutation $(i, k)(j, l)$. Then, $c(\pi) = n - 2$, $S_T(\pi) = 8$, and hence, $f_T(\pi) = 6$, but $\text{dist}_\Gamma(I, \pi) \leq 4$, as can be easily verified by applying transpositions $(j, k), (i, j), (k, l)$ and (j, k). □

Corollary 2.2. *Let T be the star $K_{1,n-1}$ on n vertices. Then the previously known upper bound inequalities*

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

and

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

hold with equality.

Theorem 2.1 implies that if T is the path graph (which is not a star for $n \geq 4$), then there exists a $\pi \in S_n$ for which the distance upper bound is strict:

Corollary 2.3. *Let T be the path graph on n vertices. Then there exists a $\pi \in S_n$ for which*

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

Despite such a result, when taking the maximum over both sides, we obtain equality:

Theorem 2.4. *Let Γ be the Cayley graph generated by a transposition tree T. Then the diameter upper bound inequality*

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

holds with equality if T is a path.

Proof. Let T be the path graph on n vertices. It is known that $\operatorname{diam}(\Gamma) = \binom{n}{2}$ [3]. Hence, it suffices to prove

$$\max_{\pi \in S_n} f_T(\pi) = \binom{n}{2}.$$

Let $\sigma = [n, n-2, \ldots, 2, 1]$. It can be verified that $f_T(\sigma)$ evaluates to $\binom{n}{2}$. Thus, it remains to prove the bound

$$f_T(\pi) \le \binom{n}{2}, \quad \forall\, \pi \in S_n.$$

We prove this by induction on n. The assertion can be easily verified for small values of n. So fix n, and now assume the assertion holds for smaller values of n. Write π as $\pi_1 \pi_2 \ldots \pi_s$. Thus π is a product of s disjoint cycles, and suppose π_1 is the cycle that contains n. We consider three cases, depending on whether π fixes n, whether π maps n to 1, or whether π maps n to some $j \neq 1$:

(a) Suppose $\pi_1 = (n)$. Define $\pi' := \pi_2 \ldots \pi_s \in S_{n-1}$. Let T' be the tree on vertex set $[n-1]$. We have that $f_T(\pi) = c(\pi) - n + S_T(\pi) = c(\pi') + 1 - n + S_{T'}(\pi') = c(\pi') - (n-1) + S_{T'}(\pi')$, which is at most $\binom{n-1}{2}$ by the inductive hypothesis.

(b) Suppose π maps n to 1. There are a few subcases, depending on the length of π_1:

(b.1) Suppose $\pi_1 = (n, 1)$, a transposition. Define $\pi' := \pi(n,1) = (1)(n)\pi_2 \ldots \pi_s$. Then, $c(\pi') = c(\pi) + 1$ and $S_T(\pi') = S_T(\pi) - 2(n-1)$. Hence, $f_T(\pi) = c(\pi) - n + S_T(\pi) = c(\pi') - 1 - n + S_T(\pi') + 2(n-1)$. Let T'' denote the tree on vertex set $\{2, 3, \ldots, n-1\}$, and let $\pi'' = \pi_2 \ldots \pi_s$ be a permutation of the vertices of T''. Then $f_T(\pi) = 2 + c(\pi'') - 1 - n + S_T(\pi') + 2(n-1) = c(\pi'') - (n-2) + S_{T''}(\pi'') + 2(n-1) - 1$. Relabeling the vertices of T'' and the elements of π'' from $\{2, \ldots, n-1\}$ to $[n-2]$ does not change $c(\pi'') - (n-2) + S_{T''}(\pi'')$, to which we can apply the inductive hypothesis. The bound then follows.

(b.2) Suppose $\pi_1 = (n, 1, j)$, where $2 \le j \le n-1$. Define $\pi' = (n,1)(j)\pi_2 \ldots \pi_s$. It can be verified that $S_T(\pi) = S_T(\pi')$, and $c(\pi') = c(\pi) + 1$. Hence, $f_T(\pi) = f_T(\pi') - 1 \le \binom{n}{2} - 1$ by the earlier subcase (b.1). Note that $S_T(\pi) = S_T(\pi^{-1})$ and $c(\pi) = c(\pi^{-1})$, so that the bound evaluates to the same value when $\pi_1 = (n, j, 1)$ and when $\pi_1 = (1, n, j)$.

(b.3) Suppose $\pi_1 = (n, 1, j_1, \ldots, j_\ell)$ contains at least 4 elements. Define $\pi' = (n, 1)(j_1, \ldots, j_\ell)\pi_2 \ldots \pi_s$, where the first cycle of π has now been broken down into two disjoint cycles to obtain π'. Define $x := |j_1 - j_2| + \ldots + |j_{\ell-1} - j_\ell|$. Then, $S_T(\pi') = 2(n-1) + x + |j_\ell - j_1| + d$ for some d, where d is the sum of distances obtained from the remaining cycles $\pi_2, \ldots, \pi_s$. Also, $S_T(\pi) = n - 1 + |1 - j_1| + x + |j_\ell - n| + d$. Also, $c(\pi) = c(\pi') - 1$. Using the equations obtained here and substituting, we get that $f_T(\pi) = c(\pi) - n + S_T(\pi) = c(\pi') - n + S_T(\pi') - n + |1 - j_1| + |j_\ell - n| - |j_\ell - j_1|$. Using the bound $c(\pi') - n + S_T(\pi') \leq \binom{n}{2}$ of subcase (b.1), and using the fact that $|1 - j_1| + |j_\ell - n| - |j_\ell - j_1| \leq n - 1$, we get the desired bound $f_T(\pi) \leq \binom{n}{2}$.

(c) We consider two subcases. In the first subcase, 1 and n are in the same cycle of π, and in the second subcase 1 and n are in different cycles of π.

(c.1) Let $\pi = (n, j_1, \ldots, j_\ell, 1, k_1, \ldots, k_t)\pi_2 \ldots \pi_s$.

Define $\pi' = (n, j_1, \ldots, j_\ell, 1)\ (k_1, \ldots, k_t)\pi_2 \ldots \pi_s$. We show that $f_T(\pi) \leq f_T(\pi')$. This latter quantity is bounded from above by $\binom{n}{2}$ due to the earlier subcases.

(c.1.1) The subcase $\ell = 0$ has been addressed in subcases (b.2) and (b.3). Since there is a vertex automorphism of the path graph T that maps 1 to n and n to 1, the subcase $t = 0$ has also been addressed by the subcase $\ell = 0$.

(c.1.2) Now suppose $\ell = 1$ and $t = 1$. The sum of distances of elements in the cycle $(n, j_1, 1, k_1)$ and $(n, j_1, 1)(k_1)$ are equal, and $c(\pi) < c(\pi')$. Hence, $f_T(\pi) < f_T(\pi')$.

(c.1.3) Now suppose $t = 1$ and $l \geq 2$; note that by symmetry, this subcase also addresses the subcase $\ell = 1$ and $t \geq 2$. A calculation of the sum of distances $S_T(\pi)$ and $S_T(\pi')$ yields, again, that $S_T(\pi) = S_T(\pi')$. Since $c(\pi) < c(\pi')$, $f_T(\pi) < f_T(\pi')$

(c.1.4) Finally, suppose $t \geq 2$ and $\ell \geq 2$. Recall that π contains the cycle $(n, j_1, \ldots, j_\ell, 1, k_1, \ldots, k_t)$, and π' contains the two cycles $(n, j_1, \ldots, j_\ell, 1)$ and $(k_1, \ldots, k_t)$. A summation of distances due to elements in these cycles yields that $f_T(\pi) \leq f_T(\pi')$ if and only if $k_1 - k_t \leq |k_1 - k_t| + 1$, which is clearly true.

(c.2) Let $\pi = (n, j_1, \ldots, j_\ell)(1, k_1, \ldots, k_t)\pi_3 \ldots \pi_s$.

(c.2.1) Suppose $\ell = 1$, i.e. $\pi_1 = (n, j_1)$ is a cycle of π. Define a new permutation $\pi' = (n, 1)\pi'_2 \ldots \pi'_s$ that has the same type as π but with the labels of 1 and j_1 interchanged, i.e. $\pi' = (1, j_1)\ \pi\ (1, j_1)$. Then $c(\pi') = c(\pi)$. Note that $S_T(\pi)$ contains terms $|n - j_1|$ and $|j_1 - n|$, corresponding to the cycle π_1. When going from π to π', the sum of two terms of S_T is increased by an amount equal to $2|j_1 - 1|$ because the cycle (n, j_1) is replaced by the cycle $(n, 1)$. When going from π to π', the cycle containing the element 1 is now replaced by a cycle containing the element j_1, and this could contribute to a decrease in S_T by at most $2|j_1 - 1|$. Hence, $f_T(\pi) \leq f_T(\pi')$. The bound then follows from applying the earlier subcase (b.1) to π'. This resolves the case $\ell = 1$, and by symmetry, also the case $t = 1$.

(c.2.2) So now assume $\ell \geq 2$ and $t \geq 2$.

Let $\pi = (n, j_1, \ldots, j_\ell)(1, k_1, \ldots, k_t)\pi_3 \ldots \pi_s$, and

let $\pi' = (n, 1)(j_1, \ldots, j_\ell, k_1, \ldots, k_t)\ \pi_3 \ldots \pi_s$. A computation of the sum of distances in $S_T(\pi)$ and $S_T(\pi')$ yields that $f_T(\pi) \leq f_T(\pi')$ if and only if $k_1 - j_\ell + k_t - j_1 \leq |k_1 - j_\ell| + |k_t - j_1| + 1$, which is clearly true. □

3 On the AK algorithm

We describe some properties of the AK algorithm here.

Theorem 3.1. *If T is a star or a path, then the AK algorithm sorts any permutation using the minimum number of transpositions.*

Proof. Let T be the path graph, with the vertices labeled consecutively from 1 to n. Let $\pi \in S_n$ be a given permutation. Then, the AK algorithm chooses, during each step, an admissible edge $(i, i+1)$ of type A or type B. If the edge is of type A, then the marker $\pi(i)$ at vertex i reduces its distance to vertex $\pi(i)$ if the transposition $(i, i+1)$ is applied, and similarly for the marker $\pi(i+1)$ at vertex $i+1$. Hence, by the chosen labeling of the vertices, $\pi(i) > \pi(i+1)$. Thus, applying $(i, i+1)$ reduces the number of inversions of the given permutation by 1. Similarly, if the edge is of type B, applying $(i, i+1)$ reduces again the number of inversions of the given permutation by 1. Thus, in either case, after $(i, i+1)$ is applied to π, we get a new permutation which has exactly one fewer inversion than π. Thus, the AK algorithm uses exactly $\mathrm{inv}(\pi)$ transpositions to home all the markers, and it is a well-known result that this is the minimum number $\mathrm{dist}_\Gamma(I, \pi)$ of transpositions possible.

Let T be the star. The different cases in the proof of Theorem 2.1 were (a),(b),(c) and (d.1) to (d.4). Each time a transposition is applied by the AK algorithm, it picks an admissible edge of type A or type B. If the edge is of type A, then we are in case (d.1) or (d.3), in which case $f_T(\pi)$ surely reduces by 1. If the edge is of type B, then we are in case (b), in which case $f_T(\pi)$ again surely reduces by 1. Thus, the AK algorithm sorts π using exactly $f_T(\pi)$ transpositions of T and this is the minimum possible number of transpositions by Theorem 2.1. □

Theorem 3.2. *There exist transposition trees for which the diameter upper bound is strict. There exist transposition trees for which the AK algorithm uses more than the minimum number of transpositions required.*

Proof. Let T be the transposition tree on 5 vertices consisting of the 4 transpositions $(1,2),(2,3),(1,4)$ and $(1,5)$. Let $\pi = (2,4)(3,5) \in S_5$. Then $f_T(\pi) = 8$, whereas a quick simulation using GAP [8] confirms that the diameter of the Cayley graph generated by T is 7. Hence, there exist transposition trees for which the diameter upper bound inequality is strict.

Next, suppose T is the transposition tree on 7 vertices consisting of the 6 transpositions $(1,2),(2,3),(1,4)$, $(4,5),(1,6)$ and $(6,7)$. Let $\pi = (2,4)(3,5)(5,7) \in S_7$. Then the following 15 edges of T, when applied in the order given, are all admissible edges (of type A or type B), and can be used to sort π on T: $(1,2),(1,4),(1,2),(2,3),(1,2),(1,6),(6,7),(1,2),(2,3),(4,5),$ $(1,4),(1,6)$, $(6,7),(1,4),(4,5)$. However, it can be verified (with the help of a computer) that the diameter of the Cayley graph generated by T is 14. Hence, the AK algorithm can take more than the minimum required number of transpositions to sort a given permutation. □

4 Further remarks

The upper bound from [1] for the diameter of Cayley graphs generated by transposition trees was studied in this work. This formula bounds the diameter of the Cayley graph on $n!$ vertices in terms of parameters of the underlying transposition tree. We showed above that this bound is sharp for all trees of minimum diameter and for all trees of maximum diameter, but can be strict for trees that are not extremal. We also showed that for some families of graphs the AK algorithm is optimal.

Let $s(n)$ denote the number of non-isomorphic trees on n vertices and let $h(n)$ denote the number of nonisomorphic trees on n vertices for which the diameter upper bound is sharp. Then, the following table seems to result from our preliminary calculations:

Table 1: Number of trees for which the diameter upper bound is sharp

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

An open problem is to characterize (all) the remaining families of trees for which the diameter upper bound is sharp. Another direction is to investigate further properties of the AK algorithm.

References

[1] S. B. Akers and B. Krishnamurthy, A group-theoretic model for symmetric interconnection networks, *IEEE Trans. Comput.*, **38**(4) (1989), 555–566.

[2] F. Annexstein, M. Baumslag, and A. L. Rosenberg, Group action graphs and parallel architectures, *SIAM J. Comput.*, **19**(3) (1990), 544–569.

[3] C. Berge, *Principles of Combinatorics*, Academic Press, New York, 1971.

[4] J. C. Bermond, T. Kodate, and S. Perennes, Gossiping in Cayley graphs by packets, In Proceedings of the Franco-Japanese conference Brest July 95, *Lecture Notes in Computer Science*, Springer-Verlag, **1120** (1996), 301–315. 1996.

[5] N. L. Biggs, *Algebraic Graph Theory*, 2nd Edition, Cambridge University Press, Cambridge, 1994.

[6] N. L. Biggs, *Discrete Mathematics*, 2nd Edition, Oxford University Press, 2003.

[7] B. Bollobás, *Modern Graph Theory*, Graduate Texts in Mathematics vol. 184, Springer, New York, 1998.

[8] The GAP Group, *GAP - Groups, Algorithms, and Programming, Version 4.4.12*, `http://www.gap-system.org`. 2008.

[9] G. Hahn and G. Sabidussi (eds.), *Graph Symmetry: Algebraic Methods and Applications*, Kluwer Academic Publishers, Dordrecht, 1997.

[10] M. C. Heydemann, Cayley graphs and interconnection networks, In *Graph symmetry: algebraic methods and applications*, Kluwer Academic Publishers, Dordrecht, 1997, 167–226.

[11] J. H. Smith, Factoring, into edge transpositions of a tree, permutations fixing a terminal vertex, *J. Combin. Theory Ser. A*, **85** (1999), 92–95.

[12] T. P. Vaughan, Bounds for the rank of a permutation on a tree, *J. Combin. Math. Combin. Comput.*, **19** (1991), 65–81.

[13] T. P. Vaughan and F. J. Portier, An algorithm for the factorization of permutations on a tree, *J. Combin. Math. Combin. Comput.*, **18** (1995), 11–31.