An Upper Bound on the Number of Circular Transpositions to Sort a Permutation

Anke van Zuylen*, James Bieron, Frans Schalekamp, Gexin Yu[1]

Department of Mathematics, The College of William and Mary, Williamsburg, VA, 23185, USA

Abstract

We consider the problem of upper bounding the number of circular transpositions needed to sort a permutation. It is well known that any permutation can be sorted using at most $\frac{n(n-1)}{2}$ adjacent transpositions. We show that, if we allow all adjacent transpositions, as well as the transposition that interchanges the element in position 1 with the element in the last position, then the number of transpositions needed is at most $\lfloor \frac{n^2}{4} \rfloor$.

Key words: Permutations; Cayley graphs; diameter; sorting; modified bubble-sort graph

1. Introduction

We consider the diameter of the Cayley graph of a permutation group with all cyclically adjacent transpositions as generating set. For a given n, the Cayley graph has $n!$ nodes corresponding to the permutations of $\{1, \ldots, n\}$ and edges between nodes representing π and π' if π' can be obtained from π by interchanging two elements in cyclically adjacent positions.

In their seminal paper Akers and Krishnamurthy [1] initiate the study of Cayley graphs in the design of interconnection networks, where the diameter of the graph is an important measure of the quality of these networks. The Cayley graph generated by cyclically adjacent transpositions is among the Cayley graphs studied in a survey on interconnection networks by Laksmivarahan, Jwo and Dhrall [8], who refer to it as the modified bubble-sort graph. In their survey, the diameter of the modified bubble-sort graph is listed as being unknown. See also Heydemann [6] for another survey of Cayley graphs as interconnection networks. A second motivation for studying the diameter of certain Cayley graphs of permutation groups, is the problem of reconstructing the evolutionary history of the genome. We refer the reader to the book [5] for a comprehensive survey.

*Corresponding author

Email addresses: anke@wm.edu, jcbieron@email.wm.edu, frans@wm.edu, gyu@wm.com

[1]Research was supported in part by NSA grant H98230-12-1-0226.

Note that by symmetry, the diameter of the Cayley graph generated by (cyclically) adjacent transpositions is equal to the maximum over all permutations of the number of transpositions required to sort the permutation. When considering the Cayley graph generated by adjacent transpositions only, then the diameter is thus equal to the maximum number of operations it takes to bubble sort any permutation π. It is well known that this is equal to the number of inversions of π, and hence the diameter of this Cayley graph of a permutation group of order $n!$ is equal to $\frac{n(n-1)}{2}$.

When considering cyclically adjacent transpositions, the problem becomes remarkably more complex. Jerrum [7] gives a polynomial time algorithm for computing a minimum length sequence of cyclically adjacent transpositions to sort any given permutation. Jerrum's algorithm is quite sophisticated, and does not yield an easy expression for the maximum number of transpositions in the sequence. Feng, Chitturi and Sudborough [4] prove that $\lfloor \frac{n^2}{4} \rfloor$ is a lower bound, and they conjecture that this bound is tight.

Chen and Skiena [3] consider a more general problem of sorting a permutation using reversals of k consecutive elements. Note that a reversal for $k = 2$ is simply an adjacent transposition. Chen and Skiena give upper and lower bounds on the sorting distance between permutations and circular permutations, where the latter can be thought of as putting the permutation on a cycle, i.e. it identifies all n permutations that can be obtained by (repeatedly) moving every element one position to the right, and the last element to the first position. The upper bound obtained for circular permutations is $O(n^2/k + kn)$. Pevzner [9] and Bafna et al. [2] consider the same problem for the case when $k = 2$. We note that sorting circular permutations using adjacent transpositions is not the same as sorting a permutation using cyclically adjacent transpositions. For the first problem, the number of transpositions needed to sort $(n, 1, 2, 3, \dots, n-1)$ is 0, since the circular permutation $(n, 1, 2, 3, \dots, n-1)$ is equivalent to the identity. When considering the number of cyclically adjacent transpositions to sort the permutation $(n, 1, 2, 3, \dots, n-1)$, the answer is $n - 1$.

To the best of our knowledge, the best known upper bound on the number of cyclically adjacent transpositions to sort a permutation was $\frac{n(n-1)}{2}$ prior to our work, i.e., no better upper bound was known than for the case when the transposition of the first and last element is excluded. In this paper, we prove that $\lfloor \frac{n^2}{4} \rfloor$ is an upper bound on the number of cyclically adjacent transpositions needed to sort any permutation of length n, thus resolving the open question of Feng et al. [4].

2. Preliminaries

We now introduce the notation we will use.

Let π be a permutation of $\{1, \dots, n\}$. We will refer to $\pi(i)$ as the position of element i. If $\pi(i) = p$, we have $\pi^{-1}(p) = i$, i.e., $\pi^{-1}(p)$ gives the element that is in position p. We will sometimes write π^{-1} as the ordered sequence

$(\pi^{-1}(1), \pi^{-1}(2), \dots, \pi^{-1}(n))$. In the following, we will use i, j, k when we want to refer to an element in $\{1, \dots, n\}$ and p, q, r to refer to a position in $\{1, \dots, n\}$.

Given a permutation π and $p, q \in \{1, \dots, n\}$, applying the transposition (p, q) to π means that we "swap" the elements in positions p and q to obtain a new permutation $\tilde{\pi}$, where $\tilde{\pi}^{-1}(p) = \pi^{-1}(q), \tilde{\pi}^{-1}(q) = \pi^{-1}(p)$, and $\tilde{\pi}^{-1}(r) = \pi^{-1}(r)$ for all $r \in \{1, \dots, n\} \backslash \{p, q\}$.

We say a transposition (p, q) is adjacent if $q = p + 1$, and we say a transposition (p, q) is cyclically adjacent if $q \equiv p + 1 \pmod{n}$. For ease of exposition, we will use $(p, p+1)$ to denote a cyclically adjacent transposition (i.e., the fact that $p+1$ is taken modulo n is implicit). We will sometimes refer to a transposition $(p, p+1)$ when applied to π with $\pi^{-1}(p) = i, \pi^{-1}(p+1) = j$, as the *swap* (i, j) of elements i and j. We remark that a swap is denoted as an ordered pair, where the first element moves "in clockwise direction", i.e., from position p to $p+1$ and the second element moves in "counterclockwise direction", i.e., from position $p+1$ to p.

We say that $Q = (q_1, q_2, \dots, q_m)$ is a sequence of cyclically adjacent swaps for π, if for every $i = 1, 2, \dots, m$ we have that q_i is a cyclically adjacent swap for the permutation that results from performing the swaps $q_1, q_2, \dots, q_{i-1}$ in order on permutation π.

We say a permutation π is sorted by a sequence of transpositions, if we obtain the identity after the sequence of transpositions is applied to π. It is well known that any permutation π can be sorted by at most $\frac{n(n-1)}{2}$ adjacent transpositions.

We will show the following theorem. This answers an open question of Feng, Chitturi and Sudborough [4].

Theorem 1. *Given any permutation π of $\{1, \dots, n\}$, there exists a sequence of at most $\frac{n^2}{4}$ cyclically adjacent transpositions to sort π.*

To prove the theorem, we begin by reviewing results by Jerrum [7]. Given a sequence of cyclically adjacent transpositions that sort π, we consider the corresponding sequence of swaps of elements. For this sequence of swaps, we let $c(i, j)$ be the number of times swap (i, j) occurs minus the number of times swap (j, i) occurs. We define the net clockwise displacement for element i as $d(i) = \sum_{j \neq i} c(i, j)$. Then we have that

$$\sum_{i=1}^{n} d(i) = 0, \tag{1}$$

since $\sum_i d(i) = \sum_i \sum_{j \neq i} c(i, j) = \sum_{(i,j):c(i,j)>0} c(i, j) + \sum_{(j,i):c(j,i)<0} c(j, i) = \sum_{(i,j):c(i,j)>0} c(i, j) - \sum_{(i,j):c(i,j)>0} c(i, j) = 0$, where the penultimate equality uses the fact that $c(i, j) = -c(j, i)$. Since the sequence of transpositions sorts π, it must be the case that

$$\pi(i) + d(i) \equiv i \pmod{n}, \text{ for every } i \in \{1, \dots, n\}. \tag{2}$$

Note that a net displacement vector d is a property of a sequence of transpositions that sort π, and that, given permutation π, there are multiple vectors

d that satisfy (1) and (2). Jerrum's key result is a characterization of the net displacement vector d of a minimum length sequence of cyclically adjacent transpositions to sort π. We first show in Lemma 2 how, given a vector d that satisfies (1) and (2), we can find a sequence of cyclically adjacent transpositions that sort π and have net displacement vector d.

We then give the expression given by Jerrum for $c(i,j)$, the net number of times swap (i,j) occurs in this sequence, as a function of π and d. Finally, we give Jerrum's main result which characterizes the net displacement vector d that corresponds to the minimum length sequence of transpositions that sort π.

Lemma 2. *Given a net displacement vector d that satisfies (1) and (2) with respect to some permutation π, a sequence of cyclically adjacent swaps that sort π and has net displacements given by d is found by repeatedly swapping cyclically adjacent elements (i,j) such that $d(i) > d(j)$, and decreasing $d(i)$ by 1 and increasing $d(j)$ by 1.*

Proof: Note that cyclically adjacent elements (i,j) such that $d(i) > d(j)$ exist unless $d(i) = 0$ for every $i = 1, \dots, n$, since d satisfies $\sum_{i=1}^{n} d(i) = 0$. After executing the swap, we decrease $d(i)$ by 1 and we increase $d(j)$ by 1. Let $\tilde{d}$ be the new displacement vector. Note that the new displacement vector $\tilde{d}$ obtained after executing the swap satisfies (1) and that the new permutation $\tilde{\pi}$ and $\tilde{d}$ satisfy (2) since $d(i) - \tilde{d}(i) = 1$ and $p(i) - \tilde{p}(i) \equiv -1 \pmod{n}$, $d(j) - \tilde{d}(j) = -1$ and $\pi(j) - \tilde{\pi}(j) \equiv 1 \pmod{n}$ and for all $k \neq i,j$ we have $\tilde{d}(k) = d(k)$ and $\tilde{\pi}(k) = \pi(k)$. Hence, if this process terminates, then it will result in a sequence of cyclically adjacent swaps to sort π with net displacement vector d.

We now argue that $\sum_k (\tilde{d}(k))^2 < \sum_k (d(k))^2$, which implies that this process does indeed terminate.

Note that

$$\begin{aligned}\sum_k (\tilde{d}(k))^2 - \sum_k (d(k))^2 &= (d(i)-1)^2 + (d(j)+1)^2 - \left((d(i))^2 + (d(j))^2\right) \\ &= -2d(i) + 2d(j) + 2.\end{aligned}$$

Now, note that, since $\pi(i) = p$ and $\pi(j) \equiv p+1 \pmod{n}$, then $d(i) > d(j)$ implies that $d(i) \geq d(j)+2$, since π and d satisfy (1). Hence $-2d(i)+2d(j)+2 \leq -2$. □

We note that this lemma generalizes sorting π by using only adjacent transpositions (i.e., bubble sort) in a natural way: In that case, we take $d(i) = i - \pi(i)$, for every $i \in \{1, \dots, n\}$. Now, let $i = \pi^{-1}(p)$ and $j = \pi^{-1}(p+1)$. Then $d(i) = i - p$ and $d(j) = j - (p+1)$, and hence $d(i) > d(j)$, implies that $i > j - 1$, and hence, $i > j$, since $i \neq j$. So in this case, the algorithm in Lemma 2 is simply the bubble sort algorithm, in which we swap adjacent elements (i,j) if $i > j$.

We now give two results that were shown by Jerrum. The expression we use in the next lemma gives an expression for $c(i,j)$, the net number of times swap

(i, j) occurs in the sequence resulting from Lemma 2. The expression is essentially the same as the expression derived on page 283 of [7]. For completeness, we give a proof in the appendix.

Lemma 3 (Jerrum[7]). *Given a displacement vector d that satisfies (1) and (2) with respect to some permutation π, any sequence of cyclically adjacent transpositions that sorts π and has net displacement vector d has*

$$c(i,j) = \begin{cases} 1 + \max\{m : \pi(i) + d(i) > \pi(j) + d(j) + mn\} & \text{if } \pi(i) < \pi(j), \\ \max\{m : \pi(i) + d(i) > \pi(j) + d(j) + mn\} & \text{if } \pi(i) > \pi(j). \end{cases}$$

Jerrum's main result is a characterization of the net displacement vector d of the minimum length sequence of cyclically adjacent transpositions for a given permutation π. The following theorem summarizes the results in Corollary 3.7 and Theorem 3.9 of [7].

Theorem 4 (Jerrum [7]). *A sequence of cyclically adjacent transpositions that sort permutation π is of minimum length if and only if each pair of elements is swapped at most once, and the net displacement vector d satisfies*

$$d(i) - d(j) \leq n \text{ for all } i, j \in \{1, \ldots, n\}. \tag{3}$$

We omit the proof since the results in the next section only rely on the fact that for any permutation π, there exists a sequence of cyclically adjacent transpositions that sort π and for which the net displacement vector d satisfies (3).

To find this sequence, we initialize $d(i) = i - \pi(i)$ for $i = 1, \ldots, n$. Note that $|d(i)| \leq n$ for every i.

Now, if $d(i) - d(j) > n$, then $d(i) > 0$ and $d(j) < 0$. If we subtract n from $d(i)$ and add n to $d(j)$, we obtain a new valid displacement vector d', which has $d'(i) = d(i) - n < 0$, $d'(j) = d(j) + n > 0$ and maintains $|d(k)| \leq n$ for all k. Therefore, $\sum_k |d'(k)| = \sum_k |d(k)| - |d(i)| + |d(i) - n| - |d(j)| + |d(j) + n| = \sum_k |d(k)| + 2n - 2(d(i) - d(j)) < \sum_k |d(k)|$. Hence, this process will terminate. We can then use Lemma 2 to find the corresponding sequence of cyclically adjacent transpositions.

3. An upper bound on the number of cyclically adjacent transpositions to sort π.

By the results from the previous section, we know that for any permutation π of $\{1, \ldots, n\}$, there exists a sequence of cyclically adjacent transpositions with net displacement vector d which satisfies (3) that sorts π. If we apply a cyclically adjacent swap (i, j) with $d(i) > d(j)$, then $d(i)$ is decreased by one, and $d(j)$ is increased by one. Hence, $\frac{1}{2} \sum_{i=1}^{n} |d(i)|$, where d satisfies (1), (2) and (3), is a *lower bound* on the number of cyclically adjacent transpositions needed to sort π. The maximum value this lower bound can take is $\frac{n^2}{4}$, since d satisfies (3). It was shown in [4] that this bound is tight for the permutation $\pi^{-1} = (\frac{n}{2} + 1, \frac{n}{2} + 2, \ldots, n - 1, n, 1, 2, \ldots, \frac{n}{2})$, where n is even.

One might conjecture that there always exists a swap such that $\frac{1}{2}\sum_{i=1}^{n}|d(i)|$ decreases by one, which would prove that $\frac{n^2}{4}$ is also an *upper bound* on the number of cyclically adjacent transpositions needed to sort any permutation π. However, this is only if there exists an adjacent swap (i, j) where $d(i) > 0$ and $d(j) < 0$. The following example shows that such a swap does not always exist: let $\pi^{-1} = (3, 2, 1, 4)$. Then $d = (2, 0, -2, 0)$ satisfies (1), (2) and (3), but any cyclically adjacent swap does not decrease $\sum_{i=1}^{n}|d(i)|$. In this section, we use different techniques to show that the conclusion does hold that at most $\frac{n^2}{4}$ cyclically adjacent transpositions suffice to sort any permutation π.

We begin by stating two auxiliary lemmas. To maintain the flow of the argument, we defer their proofs until later in this section. We slightly generalize the notion of permutation to be a bijection of any set S of positive integers. Note that the net displacement vector of a sequence of cyclically adjacent swaps is still well defined as $d(i) = \sum_{j\neq i} c(i, j)$. The following lemma follows from Lemma 3.

Lemma 5. *Let π be a permutation of S, and let $Q = (q_1, q_2, \ldots, q_m)$ be a sequence of cyclically adjacent swaps for π, with net displacement vector d that satisfies (3). Then for any two distinct elements $i, j \in S$,*

(a) $d(i) \geq d(j)$ implies $0 \leq c(i, j) \leq 1$.
(b) $d(i) = d(j)$ implies $c(i, j) = 0$.
(c) $d(i) - d(j) = n$ implies $c(i, j) = 1$.

We will prove our main result by induction, and in order to use the inductive hypothesis, we will remove some element k from S. We now define what we mean by a permutation corresponding to π restricted to $S \setminus \{k\}$. First of all, we define the relationship "i is directly before j in a permutation π" if either $\pi(i) < \pi(j)$ and there is no $\ell \in S$ so that $\pi(i) < \pi(\ell) < \pi(j)$, or $\pi(i) = \max S$ and $\pi(j) = \min S$, where $\max S$ and $\min S$ are the largest and smallest integer in S respectively. We denote this relationship by $i \prec j$.

Given a permutation π of S and a permutation $\pi' : S \setminus \{k\} \to S \setminus \{k\}$. We say π' is a permutation corresponding to π restricted to $S \setminus \{k\}$ if π' preserves the $\prec$-relationship, i.e. $i \prec j$ in π' for $i, j \in S'$ if $i \prec j$ in π, or $i \prec k \prec j$ in π. We note that there are $|S| - 1$ distinct permutations corresponding to π restricted to $S \setminus \{k\}$, but this will not be important for our purposes.

Lemma 6. *Let π be a permutation of S, and let $Q = (q_1, q_2, \ldots, q_m)$ be a sequence of cyclically adjacent swaps for π, resulting in permutation σ. Let Q' be the sequence of swaps, where all swaps involving element k are deleted, and let π' be any permutation corresponding to π restricted to $S \setminus \{k\}$. Then Q' is a sequence of cyclically adjacent swaps for π', and performing Q' on π' results in a permutation corresponding to σ restricted to $S \setminus \{k\}$.*

We now rephrase our main theorem, and use the two previous lemmas to prove it. By the results of the previous section, we have that for any permutation π there exists a minimum length sequence of cyclically adjacent transpositions that sorts π, so that every pair of elements is swapped at most once, and the net displacement vector satisfies (3). It is therefore sufficient to prove that a sequence of cyclically adjacent swaps with the properties that each pair of elements is swapped at most once, and the net displacement vector satisfies (3), has length at most $\frac{n^2}{4}$, where n is the number of elements.

Theorem 7. *Consider a sequence of cyclically adjacent swaps for a permutation π of a set of n elements, where each pair of elements is swapped at most once, and for which the net displacement vector d satisfies (3). Then the sequence consists of at most $\frac{n^2}{4}$ swaps.*

Proof: We prove the lemma by induction on n. If $n = 2$, then the lemma is clearly true, as in this case there is only one pair of elements, and this pair can be swapped at most once.

Now, assume the lemma is true for $n' = n - 1$. Consider a sequence of cyclically adjacent swaps for a permutation π of a set S of n elements that satisfies the conditions of the lemma.

Let $d_{\max} = \max_i d(i)$ and let $d_{\min} = \min_i d(i)$. Note that $d_{\max} - d_{\min} \leq n$ by (3), and hence either $d_{\max} \leq \frac{n}{2}$ or $-d_{\min} \leq \frac{n}{2}$. In the first case, let k be an element such that $d(k) = d_{\max}$; in the second case, we let k be such that $d(k) = d_{\min}$.

In order to use the inductive hypothesis, we remove element k from S. We let Q' be the sequence of swaps, where all swaps involving element k are deleted, and we let π' be any permutation corresponding to π restricted to $S \setminus \{k\}$. By Lemma 6, Q' is a sequence of cyclically adjacent swaps for π'. For the sequence Q', let $d'(i)$ be the net clockwise displacement of element i for any $i \in S \setminus \{k\}$. Note that $d'(i) = d(i) - c(i, k)$. Below, we will show that for any $i, j \in S \setminus \{k\}$, we have $d'(i) - d'(j) \leq n - 1$. Hence, Q' corresponds to a sequence of cyclically adjacent swaps for a permutation of $S \setminus \{k\}$ with net displacement vector satisfying (3) in which each pair of elements is swapped at most once. By the inductive hypothesis, Q' can have at most $\frac{(n-1)^2}{4}$ swaps. In addition, we will show that k is involved in exactly $|d(k)|$ swaps. Since $|d(k)| \leq \frac{n}{2}$, we conclude that the total number of swaps in the original sequence is at most $\frac{(n-1)^2}{4} + \frac{n}{2} = \frac{n^2+1}{4}$. Since the number of swaps is integer, it can thus be at most $\lfloor \frac{n^2+1}{4} \rfloor = \frac{n^2}{4}$.

To prove the two claims, we use Lemma 5. First, suppose that $d(k) = d_{\max}$. Then

$$\begin{aligned} d'(i) - d'(j) &= d(i) - c(i,k) - (d(j) - c(j,k)) \\ &= d(i) + c(k,i) - d(j) - c(k,j) \\ &\leq d(k) - d(j) - c(k,j) \\ &\leq n - 1. \end{aligned}$$

The first inequality uses the fact that $d(k) \geq d(i)$, so that $c(k,i) \leq 1$ by property (a) of Lemma 5, and $c(k,i) = 0$ if $d(k) = d(i)$ by property (b). The second inequality uses the fact that $d(k) - d(j) \leq n$ by (3), plus the fact that $c(k,j) \geq 0$ by property (a), and $c(k,j) = 1$ if $d(k) - d(j) = n$ by property (c).

The proof when $d(k) = d_{\min}$ is similar, and is included for completeness. In this case, we write

$$\begin{aligned} d'(i) - d'(j) &= d(i) - c(i,k) - (d(j) - c(j,k)) \\ &\leq d(i) - c(i,k) - d(k) \\ &\leq n - 1. \end{aligned}$$

The first inequality uses the fact that $-d(j) \leq -d(k)$, by definition of k, and property (a) and (b) in Lemma 5. The second inequality uses the fact that $d(i) - d(k) \leq n$ by (3), plus property (a) and (c) from Lemma 5.

Finally, note that the number of swaps in which k is involved is $\sum_i |c(i,k)|$, and by (a) and the definition of k, this is exactly equal to $|d(k)|$. □

We conclude by giving the proofs of Lemma 5 and Lemma 6.

Proof of Lemma 5:

Let τ be the permutation of S obtained by applying Q to π. Relabel the elements of S with $\{1, \ldots, n\}$ so that τ is equal to the identity, and use the same relabeling on π and Q. Then, π is a permutation of $\{1, \ldots, n\}$ and Q sorts π. Hence, the net displacement vector d corresponding to Q and π satisfy the conditions of Lemma 3. For any (i, j) we thus have

$$c(i,j) = \begin{cases} 1 + \max\{m : \pi(i) + d(i) > \pi(j) + d(j) + mn\} & \text{if } \pi(i) < \pi(j), \\ \max\{m : \pi(i) + d(i) > \pi(j) + d(j) + mn\} & \text{if } \pi(i) > \pi(j). \end{cases}$$

We take i, j such that $d(i) \geq d(j)$. Note that (3) implies that $d(j) \leq d(i) \leq d(j) + n$.

Suppose that $\pi(i) > \pi(j)$. Note that $\pi(i) < \pi(j) + n$, since π is a permutation. Hence

$$\pi(j) + d(j) < \pi(i) + d(i) < \pi(j) + d(j) + 2n,$$

so $c(i,j) = \max\{m : \pi(i) + d(i) > \pi(j) + d(j) + mn\} \in \{0, 1\}$. Moreover, if $d(i) = d(j)$, then we have $\pi(j)+d(j) < \pi(i)+d(i) < \pi(j)+d(j)+n$, so $c(i,j) = 0$, and if $d(i) = d(j) + n$, then $\pi(i) + d(i) > \pi(j) + d(j) + n$, so $c(i,j) = 1$.

Similarly, if $\pi(i) < \pi(j) < \pi(i) + n$, then

$$\pi(j) + d(j) - n < \pi(i) + d(i) < \pi(j) + d(j) + n.$$

Since $c(i,j) = 1 + \max\{m : \pi(i) + d(i) > \pi(j) + d(j) + mn\}$ we get that $0 \leq c(i,j) \leq 1$. If $d(i) = d(j)$, then $\pi(j) + d(j) - n < \pi(i) + d(i) < \pi(j) + d(j)$, so $c(i,j) = 0$, and if $d(i) = d(j)+n$, then $\pi(j)+d(j) < \pi(i)+d(i)$, so $c(i,j) = 1$. □

Proof of Lemma 6:

We prove the lemma by induction on the number of swaps m. For $m = 0$ the claim is vacuously true. For general m, denote by τ the permutation that results from performing the first $m - 1$ swaps in Q on π, and τ' the permutation that results from performing the corresponding swaps in Q' on π'. By the inductive hypothesis we know that τ' is a permutation corresponding to τ restricted to $S \setminus \{k\}$. We now discern two cases.

(case 1) k is not swapped by $q_m = (i, j)$, i.e. $i \neq k$ and $j \neq k$. We note that $i \prec j$ in τ because $q_m = (i, j)$ is a valid cyclically adjacent swap for τ. Then $i \prec j$ in τ' by the fact that τ' is a restricted permutation, so q_m is also a valid cyclically adjacent swap for τ'.

Performing one cyclically adjacent swap only changes the $\prec$-relationship for pairs of elements for which at least one element is in $\{i, j\}$. Let ℓ be so that $\ell \prec i$ in τ. Then $\ell \prec j$ in the permutation that results after swapping $q_m = (i, j)$ in τ. If $\ell \neq k$ then $\ell \prec i$ in τ' as well, and therefore $\ell \prec j$ in the permutation that results after performing swap q on τ'. If $\ell = k$, then let ℓ' be so that $\ell' \prec k$ in τ, which means that $\ell' \prec i$ in τ'. After performing q_m on τ', we have $\ell' \prec j$.

Checking the $\prec$-relations in the restricted permutation for the element right after j in τ proceeds similarly.

(case 2) k is an element that is swapped by $q_m = (i, j)$, i.e. $k \in \{i, j\}$. We let τ' be any permutation corresponding to τ restricted to $S \setminus \{k\}$. Let ℓ be the element in $\{i, j\}$ that is not equal to k. Let a and b be so that $a \prec i \prec j \prec b$ in τ. Then $a \prec \ell \prec b$ in τ'. Also, we know that $a \prec j \prec i \prec b$ in the permutation σ obtained after applying q_m to τ. This means that a permutation corresponding to σ restricted to $S \setminus \{k\}$ will have $a \prec \ell \prec b$ as well. In other words, τ' is a permutation corresponding to σ restricted to $S \setminus \{k\}$, since it has the required $\prec$-relationship between the elements. □

Acknowledgement

The authors acknowledge Chi-Kwong Li for suggesting this research topic, for leading weekly research meetings, and for numerous helpful discussions.

References

[1] Sheldon B Akers and Balakrishnan Krishnamurthy. A group-theoretic model for symmetric interconnection networks. *Computers, IEEE Transactions on*, 38(4):555–566, 1989.

[2] Vineet Bafna, Donald Beaver, Martin Fürer, and Pavel A Pevzner. Circular permutations and genome shuffling. In *Comparative Genomics*, pages 199–206. Springer, 2000.

[3] Ting Chen and Steven Skiena. Sorting with fixed-length reversals. *Discrete Applied Mathematics*, 71(1-3):269–295, 1996.

[4] Xuerong Feng, Bhadrachalam Chitturi, and Hal Sudborough. Sorting circular permutations by bounded transpositions. *Advances in Computational Biology: Advances in Experimental Medicine and Biology*, Volume 680:725–736, 2010.

[5] Guillaume Fertin, Anthony Labarre, Irena Rusu, Eric Tannier, and Stphane Vialette. *Combinatorics of Genome Rearrangements*. The MIT Press, 1st edition, 2009.

[6] Marie-Claude Heydemann. Cayley graphs and interconnection networks. In Gena Hahn and Gert Sabidussi, editors, *Graph Symmetry*, volume 497 of *NATO ASI Series*, pages 167–224. Springer Netherlands, 1997.

[7] Mark Jerrum. The complexity of finding minimum-length generator sequences. *Theor. Comput. Sci.*, 36:265–289, 1985.

[8] S Lakshmivarahan, Jung-Sing Jwo, and S.K Dhall. Symmetry in interconnection networks based on cayley graphs of permutation groups: A survey. *Parallel Computing*, 19(4):361 – 407, 1993.

[9] Pavel Pevzner. *Computational molecular biology: an algorithmic approach*. MIT press, 2000.

A. Proof of Lemma 3

Proof: Consider a permutation π and a net displacement vector d that satisfies the conditions of the lemma. First, we note that c is skew symmetric, i.e., $c(i,j) = -c(j,i)$ for any $i \neq j$. To see this, suppose without loss of generality that $\pi(i) < \pi(j)$, and let $c(i,j) = m$. Then $\pi(i) + d(i) > \pi(j) + d(j) + (m-1)n$ and $\pi(i) + d(i) < p(j) + d(j) + mn$ (where the inequality is strict because of (2)). Therefore $\pi(j) + d(j) < \pi(i) + d(i) + (1-m)n$ and $\pi(j) + d(j) > \pi(i) + d(i) + (-m)n$, so $c(j,i) = -m$.

Applying a transposition $(p, p+1)$ results in a new permutation $\tilde{\pi}$ and a new displacement vector $\tilde{d}$. Let $k = \pi^{-1}(p)$ and $\ell = \pi^{-1}(p+1)$, then $\tilde{\pi}$ and $\tilde{d}$ are given by $\tilde{\pi}(i) = \pi(i), \tilde{d}(i) = d(i)$ for $i \neq k, \ell$ and $\tilde{\pi}(k)$ is $p+1$ if $p < n$ and 1 if $p = n$, $\tilde{d}(k) = d(k) - 1, \tilde{\pi}(\ell) = p, \tilde{d}(\ell) = d(\ell) + 1$. It is clear that $\tilde{d}$ satisfies the conditions of the lemma for $\tilde{\pi}$. Let

$$\tilde{c}(i,j) = \begin{cases} 1 + \max\{m : \tilde{\pi}(i) + \tilde{d}(i) > \tilde{\pi}(j) + \tilde{d}(j) + mn\} & \text{if } \tilde{\pi}(i) < \tilde{\pi}(j), \\ \max\{m : \tilde{\pi}(i) + \tilde{d}(i) > \tilde{\pi}(j) + \tilde{d}(j) + mn\} & \text{if } \tilde{\pi}(i) > \tilde{\pi}(j). \end{cases}$$

In order to prove the lemma, we need to show that swapping (k, ℓ) decreases $c(k, \ell)$ by one, i.e., $\tilde{c}(k, \ell) = c(k, \ell) - 1$. By skew symmetry, this also implies that $\tilde{c}(\ell, k) = c(\ell, k) + 1$. In addition, we need to show that for any pair $(i,j) \neq \{(k,\ell), (\ell,k)\}$, $c(i,j) = \tilde{c}(i,j)$.

We first consider a pair of elements $(i,j) \notin \{(k,\ell), (\ell,k)\}$. If $p < n$, then $\tilde{c}(i,j) = c(i,j)$ because $\pi(i) + d(i) = \tilde{\pi}(i) + \tilde{d}(i)$ for every $i \in \{1, \ldots, n\}$ and the relative order of all pairs of elements, except (k, ℓ) is the same in π and $\tilde{\pi}$. If $\pi = n$, then the relative order of every pair containing k or ℓ is changed, but it is easily verified that the fact that $\tilde{\pi}(k) + \tilde{d}(k) = \pi(k) + d(k) - n$ and $\tilde{\pi}(\ell) + \tilde{d}(\ell) = \pi(\ell) + d(\ell) + n$ implies that $c(i,j) = \tilde{c}(i,j)$ unless (i,j) is (k,ℓ) or (ℓ, k).

We now consider the pair (k,ℓ), and show that $\tilde{c}(k,\ell) = c(k,\ell) - 1$. If $p < n$, then $\pi(k) < \pi(\ell)$ and $\tilde{\pi}(k) > \tilde{\pi}(\ell)$. Also, $\tilde{\pi}(k) + \tilde{d}(k) = \pi(k) + d(k)$ and $\tilde{\pi}(\ell) + \tilde{d}(\ell) = \pi(\ell) + d(\ell)$. Hence, $\tilde{c}(k,\ell) = \max\{m : \pi(k) + d(k) > \pi(\ell) + d(\ell) + mn\}$ and $c(k,\ell) = 1 + \max\{m : \pi(k) + d(k) > \pi(\ell) + d(\ell) + mn\}$. We thus have that $\tilde{c}(k,\ell) = c(k,\ell) - 1$. If $p = n$, then $\pi(k) > \pi(\ell)$, $\tilde{\pi}(k) < \tilde{\pi}(\ell)$, $\tilde{\pi}(k)+\tilde{d}(k) = \pi(k)+d(k)-n$ and $\tilde{\pi}(\ell)+\tilde{d}(\ell) = \pi(\ell)+d(\ell)+n$. Therefore, $\tilde{c}(k,\ell) = 1+\max\{m : \pi(k)+d(k)-n > \pi(\ell)+d(\ell)+n+mn\} = -1+\max\{m' : \pi(k)+d(k) > \pi(\ell) + d(\ell) + m'n\}$ and $c(k,\ell) = \max\{m : \pi(k) + d(k) > \pi(\ell) + d(\ell) + mn\}$. Hence, we again have that $\tilde{c}(k,\ell) = c(k,\ell) - 1$. $\square$