Minimizing Transpositions in Permutations with Indistinct Variables

Mary K. Hussey
Carthage College
mhussey@carthage.edu

May 21, 2019

Abstract

In any permutation within a symmetric group, the minimum number of transpositions required to return the permutation to the identity can be algorithmically determined. In fact, this algorithm determines this required number of transpositions based on the degree of the group and the cycle structure of the given permutation. This consistent property changes as the set acted upon by S_n becomes a multiset. In this paper, we explore how introducing these duplicates of some of the variables may reduce the required number of transpositions and determine the altered algorithm for this change in required number of transpositions. We find that this reduction is reliant on the number of indistinct variables for each unique variable, as well as the distribution of variables in the permutation and the cycles within the permutation. Since there are more intricacies to this algorithm than the base case of strictly unique variables, for any given cycle structure there are multiple possibilities of the required number of transpositions to return all variables to their original position.

1 Introduction

Suppose G is a symmetric group. It is well known and has historically been proven that any element, p, representing a permutation in G, can be written as the composition of some number of transpositions. This number of transpositions has been proven to not be unique except in terms of being even or odd, i.e., the sign of the permutation. However, it can be proven that every element p has some minimum number of transpositions necessary to represent p. More precisely, this minimum number, denoted t, can be predicted from the number disjoint cycles in p and their lengths.

As we explore this proof, we discover the interesting case where two or more variables in p are essentially indistinguishable; that is, variables for which it is the case that if they were to be swapped in position, there would be no apparent change in the ordering of the variables. In this case, these aforementioned

qualities are not sufficient to predict the minimum number of transpositions required to represent p. In fact, it becomes apparent that this minimum number of transpositions can be determined by the number of indistinguishable variables that occur within a cycle in conjunction with the number of disjoint cycles and their lengths. Since the number of indistinguishable variables that occur within a cycle is a consideration of this calculation, it can be noted that two permutations, p_1 and p_2, with the same cycle structure may have two distinct numbers of required minimum transpositions for their representations.

2 Definitions and Development

Definition 1. The **symmetric group**, denoted S_n, is the set of all permutation of some set $A := \{1, 2, ..., n\}$ and we call n the degree of the group. The group operation is function composition.

For the purposes of this paper function composition will be handled left to right. That is, for example, the composition of two permutations f and g where $f := (1, 2)$ and $g := (2, 3)$ composed as $fg = (1, 2)(2, 3)$ will first transpose the 1 and 2 followed by the transposition of the 2 and 3 so that if the original order of these variables was 1, 2, 3 they would then end in the positions 2, 3, 1. It should also be noted that any use of an arbitrary transposition will refer to a transposition contained in the elements of the functioning S_n.

Definition 2. Indistinguishable variables are some two or more variables in some permutation which when transposed make no observable difference to the order of variables.

Example 3. If the variables of S_4, $\{1, 2, 3, 4\}$, were relabeled using the letters $MOON$, such that $1 = M$, $2 = O$, $3 = O$, $4 = N$. While 1 and 2 are clearly different variables, they are indistinguishable from each other and thus transposing the them would make no difference to the spelling of the word $MOON$.

We will let p be a permutation from S_n. The following notations will be used herein:

$$n := \text{the degree of the group,}$$

$$s := \text{the number of disjoint cycles in } p,$$

$$k := \text{the number of unique variables } p^1,$$

$$i := \text{some integer such that } 1 \leq i \leq k \ ,$$

$c_i :=$ the number of disjoint cycles containing the i^{th} unique variable, and

[1] For example, if S_4 is acting upon $\{1, 1, 2, 3\}$, $k = 3$.

$m_i :=$ the number of occurrences of the i^{th} unique variable in p not interrupting any other two indistinguishable variables within a cycle.

It should be noted that interrupting refers to the occurrence where within a given cycle different repeated variables occur between each other. For example, in (a, b, a, b) the indistinguishable variables interrupt each other but in (a, a, b, b) they do not. In the case that multiple sets of indistinct variables occur within the same cycle, the variable with the most occurrences in the cycle is deemed non-interrupted as well as any subsets of other variables which are also non-interrupted.

Lemma 4 (Action of a Transposition)**.** *Any transposition composed with any element p, such that $p \in S_n$, will change the number of transpositions remaining to return the n variables to their original position by at most 1.*

Proof. We will prove that any transposition $t := (a, b)$ when composed with some permutation p can change the number of transpositions[2] it takes to return all of the variables shuffled by p back to their starting positions by ± 1. We will prove this by cases, where Case I is when a and b are in the same cycle in p and Case II is where a and b are in two disjoint cycles in p.

Case I. We will begin by assuming p's cycle structure is a single n-cycle containing all variables in the sole cycle. Without loss of generality in means of finding some number of transpositions to return all the variables of p to their original position[3], we will be using the common algorithm to find this number. That is, for some permutation $(1, 2, 3, ..., x)$ we can return all variables to their original position by composing the permutation with $(1, 2)(1, 3)...(1, x)$ where the total number of transpositions used is equal to $x - 1$ because we transpose 1 with all other $x - 1$ variables, we can determine that it would take $n - 1$ transpositions to return all the variables in p to the identity.

If we were to then compose p with some transposition (a, b), we end up with a new permutation, p_1, where we have split up the sole n-cycle into 2 cycles of size q and $n - q$. Using the same common algorithm to count the number of transpositions used to return p to the identity, we know we can count the 2 new cycles separately because they are disjoint. So, we can say the q-cycle uses $q - 1$ transpositions to return to the identity, while the $(n - q)$-cycle uses $(n - q) - 1$ transpositions to return to the identity. In total, we know p_1 uses $(q - 1) + ((n - q) - 1)$, or $n - 2$, transpositions to return to the identity. Thus, if p is composed of one n-cycle any transposition, t, will only ever change the number of transpositions used to return all of the variables in p to their original positions by 1.

If we expand the case so that p is made up of more than one cycle, but a and b are still contained within one cycle, this conclusion holds as the reduction of one permutation as proven above holds for the cycle containing a and b and the remainder of the transpositions required would be counted as normal.

[2]This is, of course, assuming the method of counting said transpositions remains constant.
[3]In other words, return p to the identity.

Case II. In this case, we will assume p has a q-cycle and an $(n-q)$-cycle, where $0 < q < n$ and transposition t transposes some two elements, (a, b), where a and b are in separate cycles. Using the above detailed common algorithm for counting transpositions we can return p to the identity with $((q-1)+(n-q)-1)$, or $n-2$, transpositions where $q-1$ is the number of transpositions used to return the q-cycle to the identity and $(n-q)-1$ is the number of transpositions used to return the $(n-q)$-cycle to the identity. If we were to compose p with t, we would join the two disjoint cycles into one n-cycle. By the common algorithm for counting transpositions, we know that this n-cycle would use $n-1$ transpositions to return to the identity. Thus if p is composed of 2 disjoint cycles, a transposition of any two elements in separate cycles can only increase the number of transpositions used to return the elements to their original positions by $+1$. If p is composed of more than 2 disjoint cycles, this holds true by the change of counting occurring in the two cycles containing variables a and b being transposed as proven above and the number of transpositions used to return the remaining cycles holds constant as the cycles are disjoint and are not affected by the transposition. Thus, any permutation, p, composed with some transposition which transposes two elements in separate cycles in p will increase the number of transpositions remaining to return p to the identity by $+1$ assuming the method of counting the transpositions remains constant.

Therefore, we have proven that, while the means of counting transpositions used to return some permutation to the identity is unchanged, the composition of any permutation p with some transposition t can only change the number of transpositions used to return the variables of p to their original position by ± 1. □

Using Lemma 4, we can prove a means of counting remaining transpositions that will always result in the minimum number of transpositions required to return the permutation to the identity in the base case where all variables in p are unique. It should also be noted that some composition of transpositions that when composed with p results in the identity, is simply a decomposition of p^{-1}. Since any transposition's inverse is itself, it is evident that the minimum number of transpositions required to return some p to the identity is also the minimum number of transpositions required to decompose p into transpositions.

Theorem 5 (Minimum Transpositions for $p \in S_n$)**.** *If all variables in p are distinct, the minimum number of transpositions required to compose p to result in the identity will be the degree of S_n minus the number of disjoint cycles*[4] *in p, or*

$$\textit{minimum transpositions} = n - s.$$

Proof. We will prove that for some permutation, p, where all of the variables are distinct (or $k = n$ such that n is the number of variables in p) the minimum number of transpositions required to compose with p to return to the identity is $n - s$. Noting that the trivial case where p only has 1 variable requires a

[4] All fixed points will be considered disjoint 1-cycles.

minimum of 0 transpositions because it the identity, we will proceed to the non-trivial cases. By Lemma 4, we know that any transposition will change the count of transpositions remaining to reach the identity by either ± 1. We can thus assume the minimum number of transpositions required will only be found when a transposition decreases the number of remaining transpositions by 1 whenever possible. To decrease the number of transpositions remaining to return to the identity, we must return either one or both of the variables being transposed to its original position. As seen in the proof of Lemma 4, a decrease of 1 transposition can only occur if the two variables are in the same cycle. Thus, we can count the minimum possible transpositions of the base case of 1 n-cycle, by noting that composing the n-cycle with some transposition can only return one of the variables to its original position unless $n \leq 2$.

If $n = 2$, the only possible transposition will return the only two variables to their original positions and the minimum number of transpositions required to compose with p to result in the identity is 1. In the case that $n > 2$, as the resulting permutation between the composition of any transposition and p will be a 1-cycle and an $(n-1)$-cycle. We know that any transposition that includes any variable in a 1-cycle will only add 1 to the number of transpositions remaining since the transposition would have to contain variables from two separate cycles. Hence, we will not be considering any variables already in their original positions as eligible to be transposed.

At this point, we can count each next required transposition that can return one of the variables to its original position in the same manner until we have $(n-2)$ 1-cycles and one 2-cycle. At which point, the final transposition will return both variables of the same cycle to their original position. If we were to count the number of transpositions used, we can look at the position just prior to the final transposition where we have $(n-2)$ 1-cycles and one 2-cycle. We have shown that each variable in its own 1-cycle at this point required a corresponding transposition or, in total $(n-2)$ transpositions and we have one transposition remaining. Thus, we require $n-1$ transpositions when some p is composed of n variables.

Using this idea that some singular cycle with n variables requires $n-1$ transpositions to return all variables in the cycle to their original position, we can count the minimum number of transpositions for a permutation containing c disjoint cycles.

It is important to note that we will now denote the number of variables in each disjoint cycle where $a \leq s$ as n_a. In this case p is composed of s cycles, so the minimum number of transpositions required to return p to the identity can be summed as

$$\sum_{a=1}^{s} (n_a - 1).$$

We know that n_a, for all values of a where $1 \leq a \leq s$, will sum to n and the -1 for each iteration of a will sum to $-s$. In other words we know the minimum number of transpositions required will be $n-s$ for any cycle structure of p where all the variables of p are distinct. Since this count of required transpositions

counts only those that decrease the remaining number of transpositions required by 1 and Lemma 4 proves that to be the only way to reduce the number of remaining required transpositions, we have proven that the minimum number of transpositions required to return some permutation p with all distinct variables to the identity must be

$$n - s.$$

□

We have now proven how to predict the basic behaviours of the common case of permutations we would see when all variables are distinct. However, what we find is that the base case of the minimum number of transpositions required to return some permutation to the identity of its symmetric group is missing a variable when an indistinguishable variables are introduced.

As we explore how the minimum number of transpositions to return a permutation p with indistinguishable variables to the identity, we find that with indistinguishable variables there is potential for extra savings when it comes to the required number of transpositions. This is because, in certain cases, if indistinct variables are returned to each others spots, the resulting permutation would still be the identity since the appropriate variable is in the appropriate position.

Take, for example, the letters in the word MO_1O_2N as our variables. If the permutation looked like $(M)(O_1, N, O_2)$ so that the current positioning of the letters looked like MO_2NO_1, the second O is not in its original position but because it is sitting where an O belongs and the two variables would be indistinguishable without the subscript we can leave it where it is an reduce the number of transpositions required to return to the identity by a whole transposition. However, we find this saving only to be present when the indistinguishable variables are in the same cycle.

3 Results

Lemma 6 (Minimum Required Number of Transpositions With a Single Indistinguishable Variable)**.** *If $p \in S_n$ is a permutation with two variables that are indistinguishable, then the minimum number of transpositions required to return the variables of p to their original positions will be*

$$n - s - 1$$

unless the indistinguishable variables occur in separate cycles in which case the minimum number of transpositions will $n - s$.

Proof. We can assume p is some permutation with a pair of indistinguishable variables. We will then divide the possible situations into two cases where the indistinguishable variables occur inside of the same cycle and the case that the indistinguishable variables do not occur inside of the same cycle.

Case I. We begin by noting that since p is a permutation and we could treat all of the variables as if they were unique and we can return p to the identity by composing it with $n - s$ transpositions. We should also note that at any point if the two indistinguishable variables are in each others' positions, we are satisfied as the position has been filled with the correct variable. Since any two indistinguishable variables contained within a single cycle are essentially interchangeable, any cycle containing two indistinguishable variables can be written equivalently as two disjoint cycles each containing one of the indistinguishable variables. For example, the cycle (a, b, a, c) could be written equivalently as $(a, b)(a, c)$. We can do this because the indistinguishable variables are interchangeable so the result of the cycle is equivalent to that of the two cycles created by splitting the original cycle just after each of the indistinguishable variables. This is even the case when the indistinguishable variables are next to each other because while they have no variables between them on one side, the first indistinguishable variable is, essentially, in a one-cycle because it is interchangeable with its pair whose position it assumes. We can then count the minimum possible number of transpositions required to return the permutations to the identity using this equivalent notation. Since we have just increased the number of cycles by one, we conclude that the number of transpositions required has become $n - (s + 1) = n - s - 1$. This is a new minimum and, in fact, it has to be the absolute minimum number of transpositions required because we have already counted by strictly decreasing the remaining number of transpositions required to reach the identity with each transposition applied. Thus, the minimum number of transpositions required to return the permutation to the identity is

$$n - s - 1.$$

Case II. We can use the same logic as in the first case to conclude that the minimum number of transpositions required to return p to the identity will be found by treating the separate cycles as disjoint (which they are, though they may not look it) and finding the minimum number of transpositions as if all variables were unique. Thus, in this case the minimum number of transpositions will be

$$n - s.$$

At this point we have proven that if two indistinguishable variables occur in the same cycle within a permutation the minimum number of transpositions required to compose that permutation will be

$$n - s - 1,$$

while if the two indistinguishable variable occur in separate cycles within a permutation the minimum number of transpositions required to compose that permutation will be

$$n - s.$$

□

From here, we can use this knowledge to introduce more of the same indistinguishable variable to prove that this pattern of saving in transpositions required to return to the identity will continue.

Lemma 7 (Minimum Transpositions With Repeating Indistinguishable Variables)**.** *If $p \in S_n$ is some permutation containing some n number of variables and k distinct variables, $k < n$, such that any of the m indistinct variables are indistinguishable from each other, then the minimum number of transpositions required to return p to the identity will be*

$$n - s - (m - c)$$

such that c is the number of disjoint cycles containing any occurrence of the indistinguishable variables.

Proof. Using the same logic as in the proof for Lemma 6, we can note that between any two indistinguishable variables that occur within the same cycle there exists a disjoint cycle. Thus for every occurrence of the variable beyond the first within a single cycle, there will be another disjoint cycle which by the same logic as Lemma 6 will save one transposition upon the return to the identity. Therefore, the savings of required transpositions from the base case of all distinguishable variables will be the number of occurrences of the variables minus the number of cycles they occur in as each of those cycles counts the initial indistinguishable variable in the cycle and the rest remain as the total savings of transpositions. Therefore, we say the minimum number of required transpositions to return p to the identity when p contains some m number of indistinguishable variables is

$$n - s - (m - c).$$

□

Theorem 8 (Minimum Number of Transpositions)**.** *If p is a permutation with any number of variables that are repeated*[5]*, then the minimum number of transpositions required to return the variables of p to their original positions can be found by*

$$n - s - \sum_{i=1}^{k} (m_i - c_i).$$

Proof. Using Lemma 7, we know that for any one repeated variable the permutation can be rewritten into disjoint cycles where the repeated variables end up in a different position than the original permutation results but is equivalent because of their indistinguishable nature. By this same logic, we find any nested repeated variables distinct from the outer repeated variable will create another

[5]This theorem holds for the case that none of the sets of repeated variables interrupt each other. In the case that any of the sets of indistinguishable variables are interrupted, the minimum number of transpositions will be found by not including the occurrences of the interrupting variable or variables in the count of m_i.

disjoint cycle between their repetitions. This will only be the case when the nested indistinguishable variables all lie within the outer indistinguishable variables as they will not be separated when the disjoint cycles of the outer variables are taken into account. In the event that different sets of indistinguishable variables interrupt each others' separate cycles, the number of transpositions saved by reorganizing the cycle to separate disjoint cycles will be maximized by maximizing the number of separate cycles. This means the indistinguishable variable that becomes the outer variable should always be the variable with more occurrences within each given cycle and in the case of the number of occurrences being equal either can be chosen and will accurately represent the minimum number of transpositions required to return the permutation to the identity. This will be found by taking the number of variables of p minus the number of cycles in p minus the summation of each unique variable's number of cycles containing the unique variable subtracted from the number of occurrences of that variable where the repeated occurrences lie within the bounds of any other two indistinguishable variable or are the only indistinguishable variables in the cycle or

$$n - s - \sum_{i=1}^{k} (m_i - c_i).$$

□

4 Conclusion and Directions for Further Research

This finding of potential savings for the minimum number of transpositions required to represent a permutation or return it to the identity is an exciting mathematical result in and of itself. It is certainly an interesting case when not all of the variables inside a permutation are distinguishable from each other and these findings may be applicable in some use of a bubble sort while programming. It could also be applicable to functionality of different puzzles, like word jumbles, where the goal is some kind of sorting of elements that may or may not be distinct.

To focus more mathematically, there is also potential for further research. One thing to be considered is how the sign of any given permutation may be affected since the number of transpositions required to represent the permutation may no longer be even/odd as originally thought. The challenge in such research is to determine how likely that sign change is for any given permutation based on the cycle structure and ordering of variables. If this change is not symmetrical between even and odd permutations, we could also be asking what happens to the normal subgroup in such cases.

Another route to explore would be the likelihood of this "maximum" savings of transpositions because within a permutation, depending on the cycle structure and composition of distinct and indistinct variables, it is not guaranteed. This raises the question: is any general case to predict such a likelihood for a given permutation or a given set of variables?

References

[1] Gallian J. A., (2013). *Contemporary Abstract Algebra,* 8^{th} ed. Boston, MA: Brooks/Cole