Permutations and the 15-Puzzle

Peter Trapa

January 21, 2004

1. introduction: the 15-puzzle

In the late 1870's the mathematical puzzlemaker Samuel Loyd introduced the now famous 15-puzzle. The game consists of a 4-by-4 grid together with 15 tiles numbered $1, 2, \ldots, 15$, and a single vacant location on the grid.

1	2	3	4
5	6	7	8
9	10	11	12
13	14	15	

A legal move consisted of sliding a numbered tile into the vacant location. From the initial configuration above, for instance, there are two legal moves: sliding the 12 down or the 15 to the right. The object of the puzzle is to use a sequence of legal moves to interchange the position of the tiles labeled 14 and 15 while leaving all other tiles unchanged.

Loyd offered a prize of $1000 (a princely sum in 1870) for the first correct solution. The puzzle swept quickly across America, then Europe, and in Loyd's own words eventually "drove the world crazy." This was exaggeration, but perhaps not too much so. The mathematician W. E. Story, not notable himself for hyperbole, wrote in an 1879 issue of *The American Journal of Mathematics*, America's then premier mathematical publication,

> The "15" puzzle for the last few weeks has been prominently before the American public, and may safely be said to have engaged the attention of nine out of ten persons of both sexes and of all ages and condition of the community.

The hysteria surrounding the 15-puzzle must surely have delighted Loyd: he understood from the outset that no solution is possible! In the next few weeks we'll develop the machinery to figure out why it is impossible. This turns out to be relatively easy. What is harder is to figure out what configurations *are* possible to obtain from the original one using only legal moves. This will be our ultimate goal[1].

[1] A good reference is the excellent article, "A Modern Treatment of the 15-Puzzle," by A. F. Archer in the November 1999 issue of *The American Mathematical Monthly.* Much of this introduction was drawn from that source.

2. Parity of integers

Let's start by considering the integers $\mathbb{Z}$. One of the interesting features of $\mathbb{Z}$ is the existence of the notion of *parity*, that is evenness and oddness. Somewhere in our distant past we learned the basic rules

even plus even equals even
even plus odd equals odd
odd plus odd equals even,

and from those we can quickly derive the analogous rules for multiplication, as well as slightly more complicated rules such as

the sum of an odd number of odd numbers is odd, $(*)$

and so on.

This is so obvious and natural that we hardly ever think about it. But the notion of parity can be surprisingly useful. Consider the following problem[2]. Suppose there are 33 people at a party. Then we claim that it follows that at least one person at the party knows an even number of people. (Here we assume that acquaintance is mutual — if you know someone, then they know you — and we also allow for the possibility of total strangers attending who don't know anyone at all. In the latter case, the problem is trivial since zero is an even number.) Here is how to prove our claim. For convenience, label the people at the party as $1, 2, \ldots, 31$. Let n_i denote the number of acquaintances that the ith person has. Since acquaintance is mutual, the sum $n_1 + n_2 + \cdots + n_{31}$ must equal *twice* the total number of pairs of acquaintances. Thus the sum $n_1 + n_2 + \cdots + n_{31}$ is *even.* By $(*)$ above, at least one of the n_i's must be even, and the claim is proved. In fact the same argument shows that at any party with an odd number of people, someone has an even number of acquaintances.

There are lots of other fun examples like this one. Our interest here isn't so much in discussing the notion of parity on $\mathbb{Z}$, but instead on other sets (like sets of permutations). Before we get to that, it's perhaps useful to ask if there are notions of parity on other familiar sets. For instance, can one make sense of evenness and oddness for the real numbers $\mathbb{R}$? If you think about this, you'll see that the best we can do is to consider the notion of positivity and negativity. (We could have done this for $\mathbb{Z}$ too.) Positivity behaves well for multiplications (negative times negative is positive, positive times positive is positive, etc.), but doesn't behave well with respect to addition. There simply isn't a good notion of evenness and oddness. (Even so, positivity is very useful. It turns up in applications such as showing that the integral of an odd function over an interval symmetric about the origin is identically zero, if defined.)

3. permutations

Now we turn to our objects of interest, permutations of n objects. For convenience, we might as well label these objects $1, 2, \ldots, n$. We all have a heuristic understanding of what permutations are: they're simply an arbitrary rearrangement of the entries $1, 2, \ldots, n$. (So there are exactly $n!$ of them.) But we need to develop a good formalism and notation to deal with them efficiently. We start with a definition.

[2] I found this in a book by Laslo Lovasz called *Discrete Mathematics.*

Definition. A *permutation* of n letters is a function

$$\sigma \,:\, \{1, 2, \ldots, n\} \longrightarrow \{1, 2, \ldots, n\},$$

which is one-to-one (or equivalently onto — why?). The set of all such functions is denoted S_n

With this definition in hand we can develop a good notation to work with permutations. Fix $\sigma \in S_n$. We may denote map σ as follows,

$$\sigma = \begin{bmatrix} 1 & 2 & 3 & \cdots & n \\ \sigma(1) & \sigma(2) & \sigma(3) & \cdots & \sigma(n) \end{bmatrix} \tag{3.1}$$

The condition that σ is one-to-one means that there are no repeated entries in the bottom row of (3.1); hence *the bottom row of (3.1) is a rearrangement of the numbers* $1, 2, \ldots, n$. So this notation indeed recovers our heuristic notion of a permutation.

Notice that the top row in (3.1) is really superfluous. So we typically drop it from the notation and instead write

$$\sigma = [\sigma(1)\ \sigma(2)\ \sigma(3)\ \cdots\ \sigma(n)] \tag{3.2}$$

There is another convenient notation, called the cycle notation, which is sometimes simpler in certain instances. It's best illustrated through an example. Consider

$$\sigma = [2\ 5\ 4\ 3\ 6\ 7\ 1\ 8\ 9]\,. \tag{3.3}$$

The cycle notation "follows" particular elements of the rearrangement. First we start with the element 1. According to (3.3), 1 gets mapped to 2 via σ; subsequently 2 gets mapped to 5; then 5 is mapped to 6; 6 to 7; and finally 7 gets mapped back to 1. We represent this information as a "cycle,"

$$(1\ 2\ 5\ 6\ 7)$$

meaning

$$\sigma(1) = 2;\ \sigma(2) = 5;\ \sigma(5) = 6;\ \sigma(6) = 7;\ \sigma(7) = 1. \tag{3.4}$$

Notice that there was nothing special about starting with 2; we could have started with 5 or 6 or 7. So the following cycles all represent the data of (3.4),

$$(1\ 2\ 5\ 6\ 7) = (2\ 5\ 6\ 7\ 1) = (5\ 6\ 7\ 1\ 2) = (6\ 7\ 1\ 2\ 5) = (7\ 1\ 2\ 5\ 6).$$

But σ contains more information than this single cycle represents. So we need to trace though the action of other elements. Starting with 3, we observe that 3 is mapped to 4, and then 4 gets mapped back to 3. We capture this in the cycle

$$(34) = (43).$$

Next σ maps 8 to 8, so we get a single cycle (8). Similarly we get a single cycle (9). Thus we can represent the permutation σ as a product of its cycles,

$$\sigma = (1\ 2\ 5\ 6\ 7)(3\ 4)(8)(9).$$

There are many other possibilities, arising both from the ambiguity of a single cycle and the rearrangement of multiple cycles. For example,

$$\sigma = (9)(4\ 3)(5\ 6\ 7\ 1\ 2)(3\ 4)(8).$$

Notice also once we know $\sigma \in S_9$, the extra information of the single cycles (8) and (9) is redundant; so we typically drop them from the notation and simple write

$$\sigma = (1\ 2\ 5\ 6\ 7)(3\ 4).$$

Exercise. Convert the following permutations to cycle notation.

(a) $\sigma_a = [3\ 9\ 10\ 4\ 5\ 6\ 1\ 2\ 8\ 7]$

(b) $\sigma_b = [7\ 8\ 1\ 4\ 5\ 6\ 10\ 9\ 2\ 3]$

(c) $\sigma_c = [1\ 9\ 8\ 3\ 5\ 6\ 2\ 7\ 4\ 10]$

(d) $\sigma_d = [6\ 5\ 4\ 3\ 2\ 1\ ]$

(e) $\sigma_e = [15\ 14\ 13\ 12\ 11\ 10\ 9\ 8\ 7\ 6\ 5\ 4\ 3\ 2\ 1]$

Exercise. Convert the following examples of cycle notation to row notation.

(a) $\tau_a = (2\ 5\ 6\ 8) \in S_{10}$

(b) $\tau_b = (2\ 5\ 6\ 8) \in S_8$

(c) $\tau_c = (1\ 2\ 5)(3\ 4\ 6) \in S_6$

(d) $\tau_d = (1\ 2)(2\ 3\ )(3\ 4)(4\ 5)(5\ 6) \in S_6$

The next thing to observe about elements of S_n is that we can *compose* them. That is if we have two functions

$$\sigma, \tau\ :\ \{1, 2, \ldots, n\} \longrightarrow \{1, 2, \ldots, n\},$$

we may consider the functions

$$\sigma \circ \tau, \tau \circ \sigma \; : \; \{1, 2, \ldots, n\} \longrightarrow \{1, 2, \ldots, n\}.$$

Notice that if σ and τ are one-to-one, then so are $\sigma \circ \tau$ and $\tau \circ \sigma$. *In other words, if $\sigma, \tau \in S_n$, then $\sigma \circ \tau, \tau \circ \sigma \in S_n$.* We typically drop the $\circ$ from the notation, and simply write $\sigma\tau$ and $\tau\sigma$.

Exercise. Refer to the previous two exercise and compute the following products:

(a) $\sigma_a \sigma_b$

(b) $\sigma_b \sigma_a$

(c) $\sigma_a \tau_a$

(d) $\tau_a \sigma_c$

(e) $\tau_c \tau_d$

(f) $\tau_d \tau_c$

Notice that

$$e := [1\ 2\ 3\ \cdots\ n]$$

is a very special element of S_n. It has the property that

$$\tau e = e \tau = \tau \qquad \text{for all } \tau \in S_n.$$

Moreover, it is not hard to see that for all $\tau \in S_n$, there exists $\tau^{-1} \in S_n$ such that

$$\tau \tau^{-1} = \tau^{-1} \tau = e.$$

Not surprisingly τ^{-1} is called the *inverse* of τ.

Exercise. Refer to the preceding exercise and compute the following inverses.

(a) σ_a^{-1}

(b) σ_c^{-1}

(c) σ_e^{-1}

(d) τ_a^{-1}

(e) τ_c^{-1}

(f) τ_e^{-1}

4. LENGTH AND PARITY OF PERMUTATIONS

Out next task is to introduce a notion of "complexity" of a permutation. To get an idea of what we might mean by this, consider the following three permutations in S_5:

$$\sigma_1 = [1\ 2\ 3\ 4\ 5]; \qquad \sigma_2 = [2\ 1\ 4\ 3\ 5]; \qquad \sigma_3 = [5\ 4\ 3\ 2\ 1].$$

We can visualize each by connecting the points in their two-row notation (as in (3.1)) as follows:

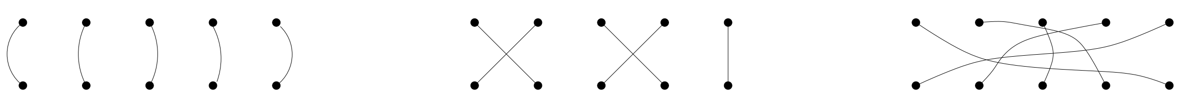

From these pictures — which we will henceforth call *string diagrams* — it seems reasonable to consider σ_1 the least complex of three. Next in complexity is σ_2. Finally σ_3 appears to be most complex. This intuition is based on the number of crossings in the string digram of each σ_i. Our goal is to make this intuition precise. We being as follows.

Provisional Definition. Define the *length* of $\sigma \in S_n$ to be the number of crossings in the string diagram of σ.

This definition is provisional because it is not precise enough. For instance, we could just as well draw the string diagrams of σ_3 as

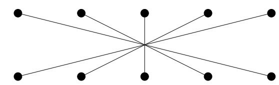

or the string diagram of σ_1 as

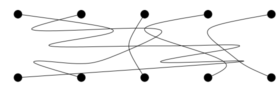

Notice that according to our first string diagram, the length of σ_3 is 10, but from the second diagram it is just 1. Similarly from the first diagram, the length of σ_1 is 0, but from the second it is definitely nonzero!

Thus the provisional definition doesn't seem to be a good one. There are a couple of ways one could fix things. The above examples show that we need to modify our notion of string diagram: we want to pull the strings taut (to rule out the case of the second picture for σ_1) and we want to spread out "multiple crossing" (to avoid the second picture for σ_3). This is a little messy to make precise, so we try a second approach and ask: When do two string cross in the refined notion of a string diagram? The answer is a little clever, but very intuitive once one gets the hang of it:

$$\text{Fix } i < j. \text{ The } i\text{th and } j\text{th string cross if } \sigma(i) > \sigma(j).$$

This is a nice precise condition, and it leads to our definition of length.

Definition. Define the length of $\sigma \in S_n$ as follows,

$$l(\sigma) = \text{ the number of pairs } (i,j) \text{ with } 1 \leq i < j \leq n \text{ and } \sigma(i) > \sigma(j).$$

For example, according to this definition,

$$l(\sigma_1) = 0; \qquad l(\sigma_2) = 2; \qquad l(\sigma_3) = 10,$$

agreeing with the crossing in the nicely drawn string diagrams we initially wrote down. In fact, it almost always useful to think of the length as counting crossing in nicely drawn string diagrams. But we have to be careful since "nicely drawn" is rather imprecise!

Exercise. Compute the length of five of the permutations appearing in the previous section.

Next we want to ask if the (revised) definition of length has any good properties. For example we can ask

$$\text{Does } l(\sigma\tau) = l(\sigma) + l(\tau)? \tag{4.1}$$

A quick moment of thought shows this can't be true. For instance, $\sigma_1 = \sigma_2\sigma_2$ but

$$l(\sigma_1) = 0 \neq 2 + 2 = l(\sigma_2) + l(\sigma_2).$$

In fact, one can easily show that $l(\sigma) = l(\sigma^{-1})$ (why?), so (4.1) fails miserably.

Can me modify the question in (4.1) to end up with something interesting?

Theorem 4.2. *The length defines a good notion of parity on S_n. More precisely, fix $\sigma, \tau \in S_n$. Then*

(1) *If $l(\sigma)$ and $l(\tau)$ are even, then $l(\sigma\tau)$ is even;*

(2) *If $l(\sigma)$ and $l(\tau)$ are odd, then $l(\sigma\tau)$ is even;*
(3) *If $l(\sigma)$ is even and $l(\tau)$ is odd, then $l(\sigma\tau)$ is odd;*
(4) *If $l(\sigma)$ is even and $l(\tau)$ is odd, then $l(\tau\sigma)$ is odd.*

(Note: since $\sigma\tau$ need not equal $\tau\sigma$, (4) need not follow from (3).)

Because of the theorem, the next definition is useful.

Definition. If $\sigma \in S_n$ and $l(\sigma)$ is even, then we call σ an *even* permutation. Similarly if $l(\sigma)$ is odd, then we say that σ is an *odd* permutation. (Thus the theorem says that the product of an even and an odd permutation is odd, and so on.)

We are now going to prove the the theorem. (You should try it on your own first!) To begin we need two lemmas.

Lemma 4.3. *For $1 \leq i \leq n-1$, define (in cycle notation),*

$$s_i = (i\ i{+}1).$$

Then any element $\sigma \in S_n$ may be written as a product of elements of the form s_i. (Note that the elements s_i are exactly the elements of length 1. So the lemma says that every permutation σ is a product of length 1 permutations.)

Proof of Lemma. The conclusion of the lemma says that any rearrangement of a list may be accomplished by interchanging adjacent entries of the list. But this is obvious. (Note that the expression in terms of the s_i need not be unique! This amounts to saying that you can accomplish the same ultimate rearrangement by swapping different sequences of adjacent entries.) $\square$

Lemma 4.4. *Fix $\sigma \in S_n$ and $1 \leq i \leq n-1$, and define s_i as in Lemma 4.3. Then either*

$$l(\sigma s_i) = l(\sigma) + 1 \qquad \text{or} \qquad l(\sigma s_i) = l(\sigma) - 1.$$

The same conclusion holds for $l(s_i\sigma)$. (In other words: multiplying σ by s_i changes its parity.)

Proof of Lemma. The proof is best visualized in terms of nicely drawn string diagrams for σ. One possibility is that the strands for the string diagram for σ which begin at i and $i+i$ cross (that is, $\sigma(i) > \sigma(i+1)$). In this case the effect of composing with s_i untangles them. No other crossings are affected. So

$$l(\sigma s_i) = l(\sigma) - 1.$$

On the other hand, the only other possibility is that strands for the string diagram for σ which begin at i and $i+i$ do not cross (that is, $\sigma(i) < \sigma(i+1)$). In this case the effect of composing with s_i tangles them. No other crossings are affected, and so

$$l(\sigma s_i) = l(\sigma) + 1.$$

The argument for $l(s_i\sigma)$ is identical. This proves the lemma. $\square$

Using the lemmas, we can now prove the theorem.

Proof of Theorem 4.2. We first prove (1). Fix $\sigma, \tau \in S_n$ and assume that both are even. Using Lemma 4.3 write

$$\sigma = s_{i_1} s_{i_2} \cdots s_{i_N}$$
$$\tau = s_{j_1} s_{j_2} \cdots s_{j_M}.$$

According to Lemma 4.4, N and M are *even* since both σ and τ are even. Now write

$$\sigma\tau = s_{i_1} s_{i_2} \cdots s_{i_N} s_{j_1} s_{j_2} \cdots s_{j_M}.$$

There are an *even* number of terms on the right side. So Lemma 4.4 says that $l(\sigma\tau)$, and hence $\sigma\tau$, is even. This completes the proof of Theorem 4.2(1). The other parts of the theorem follow in exactly the same way. $\square$

Since the even permutation have good multiplicative properties — the product of two evens is always again even — we set them aside.

Definition. The set of even permutations is called the *alternating group on n letters* and denoted A_n.

Exercises.

1. Write down the three elements of A_3 and the 12 elements of A_4.

2. Write each element of A_3 as a product of s_i's. Do the same for A_4.

3. Show that in A_4 every element is a product of 3-cycles. (Compare Lemma 4.3.)

4. Show that $s_i s_{i+1} = (i\ i{+}1\ i{+}2)$.

5. Prove that A_n has $n!/2$ elements.

Exercise (4) above suggests a general result.

Theorem 4.5. *For $1 \leq i \leq n-2$, let $t_i = (i\ i{+}1\ i{+}2) \in A_n$; so $t_i^{-1} = (i\ i{+}2\ i{+}1)$. Every element of A_n may be written as a product of 3-cycles of the form t_i and t_i^{-1}. (Compare Lemma 4.3.)*

Proof. Fix $\sigma \in S_n$. By Lemma 4.3, we may write

$$\sigma = s_{i_1} s_{i_2} \cdots s_{i_N}.$$

We know that there are an even number of terms here. So we may regroup this expression in pairs,

$$\sigma = (s_{i_1} s_{i_2})(s_{i_3} s_{i_4}) \cdots (s_{i_{N-1}} s_{i_N}). \tag{4.6}$$

It is thus enough to show that each adjacent pair of s's is a product of three cycles. So fix such a pair, say $s_i s_j$. Let us assume $i < j$. (We return to the other case in a moment.) We can write

$$s_i s_j = s_i(s_{i+1}s_{i+1})(s_{i+2}s_{i+2})\cdots(s_{j-1}s_{j-1})s_j;$$

regrouping again, we get

$$s_i s_j = (s_i s_{i+1})(s_{1+1}s_{i+2})\cdots(s_{j-1}s_j).$$

By Exercise (3) above, each pair in parenthesis is a 3-cycle; in fact, one sees that

$$s_k s_{k+1} = t_k.$$

So assuming $i < j$, each $s_i s_j$ is a product of t_i's. If on the other hand $i > j$, then the same argument would work, but $s_i s_j$ would be a product of t_k^{-1}'s. So each pair on the right-side of (4.6) is a product of t_i's or t_i^{-1}'s. So σ is too. $\square$

Consider the following three "natural" copies of A_7 inside A_{15}: let $A_7^{(1)}$ be the alternating group on $\{1, 2, \ldots, 7\}$; let $A_7^{(2)}$ be the alternating group on $\{5, 6, \ldots, 12\}$; and finally let $A_7^{(3)}$ be the alternating group on $\{10, 11, \ldots, 15\}$.

Corollary 4.7. *Every element of A_{15} may be written as a product of elements in $A_7^{(1)}$, $A_7^{(2)}$, and $A_7^{(3)}$.*

Proof. $A_7^{(1)}$ contains the elements $t_1^{\pm 1}, \ldots, t_5^{\pm 1}$; $A_7^{(2)}$ contains the elements $t_5^{\pm 1}, \ldots, t_{10}^{\pm 1}$; $A_7^{(3)}$ contains the elements $t_{10}^{\pm 1}, \ldots, t_{13}^{\pm 1}$. So the corollary follow from the theorem. $\square$

We are going to reduce the study of the 15-puzzle to the study of the individual groups $A_7^{(i)}$.

5. Warm-up: the 7-puzzle

Consider the 7-puzzle

1	2	3	4
5	6	7	

Legal moves are just like those for the 15-puzzle.

We need to define a kind of equivalence of configurations. Given an arbitrary placement of the vacant space, let us agree to migrate it to the lower right-hand corner via the following snake

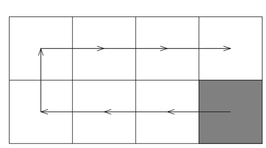

We call two configurations that differ by snaking the vacant space this way equivalent. For instance

1	2	3	4
5		6	7

and

2	3		4
1	5	6	7

are equivalent, but

1	2	3	4
5		6	7

and

2	3	4	7
1	5	6	

are not.

We need a notation for a configuration. So take an arbitrary one, snake the vacant square to the lower right corner, and assume it looks like

i_1	i_2	i_3	i_4
i_5	i_6	i_7	

We call this configuration $[i_1\ i_2 \ldots i_7]$, and may view it an as element of S_7.

First we notice that we can build the configurations $t_1 = [2\ 3\ 1\ 4\ 5\ 6\ 7] = (123)$ and $t_2 = [3\ 1\ 2\ 4\ 5\ 6\ 7] = (132)$. To see this, begin with the initial configuration and cycle the void clockwise about the 2-by-8 rectangle so as to isolate the tiles 1, 2, and 3 in the first 2-by-2 square:

1	2	3	4
5	6	7	

$\rightarrow$

1	2	3	4
	5	6	7

$\rightarrow$

	2	3	4
1	5	6	7

$\rightarrow$

2	3	4	
1	5	6	7

$\rightarrow$

2	3	4	7
1	5	6	

,

and finally

2	3	4	7
1		5	6

.

Now shuffle the void counter-clockwise in the first two-by-two square to obtain

3	1	4	7
2		5	6

.

Finally we can shuffle the void around the 2-by-8 rectangle counter-clockwise as follows

3	1	4	7
2		5	6

$\rightarrow$

3	1	4	7
2	5	6	

$\rightarrow$

3	1	4	
2	5	6	7

$\rightarrow$

	3	1	4
2	5	6	7

$\rightarrow$

2	3	1	4
	5	6	7

, ,

ultimately to arrive at

$t_1 = (1\ 2\ 3) =$

2	3	1	4
5	6	7	

.

Notice that if we had shuffled the void *twice* about the 2-by-2 square above (or if we shuffled once in the opposite direction) we would have achieved the configuration t_1^{-1}.

It is clear that by modifying our original cycle in the 2-by-8 square, we could have brought different tiles (other than 123) into the initial square. More precisely, we could have have 234, 347, 476, 567, 156, or 215. Using the same trick, we thus would build the respective configurations

$$t_2, t_2^{-1}, (3\ 4\ 7), (3\ 7\ 4), (4\ 7\ 6), (4\ 6\ 7), t_5, t_5^{-1}, (1\ 5\ 6), (1\ 6\ 5), (2\ 1\ 5), (2\ 5\ 1). \tag{5.1}$$

We want to use Theorem 4.5 (and Exercise (5) of the previous section) to conclude that by composing these moves we get all $7!/2 = 2520$ configurations corresponding to A_7. It suffices to build $t_3^{\pm 1}$ and $t_4^{\pm 1}$ as compositions from the list in (5.1). This isn't so hard to do directly; for instance,

$$t_3 = (1\ 6\ 5)(4\ 7\ 6)(3\ 4\ 6)(4\ 6\ 7)(1\ 5\ 6),$$

but one may alternatively use the following trick. If we relabel "5" by "7", then $t_1^{\pm 1}$ together with the list in (5.1) is exactly $t_1^{\pm 1}, \ldots, t_5^{\pm 1}$. So Theorem 4.5 says that using the compositions of our original list ($t_1^{\pm 1}$ and the elements of (5.1)), we get what one obtains by taking A_7 and switching "5" and "7". So compositions of our original list gives $7!/2$ elements. Since they are all even, we conclude that, indeed, compositions of our original list gives A_7.

Thus in the 7-puzzle we obtain at least the $7!/2$ configurations corresponding to A_7. But how do we know that we get no more than these? Perhaps there is a clever kind of move that we are missing.

To see that we are not missing anything, we must examine four basic moves. They correspond to taking the original configuration and sliding the void into one of the four available slots on the bottom row,

1	2	3	4
5		6	7

Next we move the void up,

1		3	4
5	2	6	7

and reshuffle the void, along our chosen snake, back to the lower right hand corner.

5	1	3	4
2	6	7	

One checks that these four moves are all even permutations. For instance, the above one gives

$$[7\ 1\ 3\ 4\ 5\ 6\ 2] = (1\ 7\ 2) = (1\ 7)(7\ 2).$$

(Of the remaining three moves, one gives the identity, and one gives the 7-cycle (1234567). You should check the remaining move!) Since there are no other possible moves (besides the inverse of these moving the void from the top row to the bottom which are also even), and since every move is a composite of these basic moves, we concluded that we can only achieve *even* configurations. We have already seen that we can achieve them all.

6. The 15-puzzle

Let's return to our original 15-puzzle. Recall Corollary 4.7. By sliding the vacant tile into one of the three adjacent pairs of rows and working in those two rows only, the discussion of the previous section shows that we may achieve the configurations corresponding to $A_7^{(1)}$, $A_7^{(2)}$, and $A_7^{(3)}$. By Corollary 4.7, we get all of A_{15}.

Can we get more? We have to consider the same kinds of basic moves of the previous section. We take a configuration, snake the void to an arbitrary location by a fixed snake, move the void up or down, and snake the void back along our fixed snake. The key is that *this operation takes place in a pair of rows.* We have done the calculation for pairs of rows already in the previous section, and so we conclude all basic moves are even. Thus only even configurations can arise, and indeed we get them all.

Finally return to Loyd's original challenge: swap the 14 and 15 while leaving all other tiles fixed. But (14 15) is an odd permutation! So this is impossible.