INVOLVE 7:1 (2014)
dx.doi.org/10.2140/involve.2014.7.41

Slide-and-swap permutation groups

Onyebuchi Ekenta, Han Gil Jang and Jacob A. Siehler

(Communicated by Joseph A. Gallian)

We present a simple tile-sliding game that can be played on any 3-regular graph, generating a permutation group on the vertices. We classify the resulting permutation groups and obtain a novel presentation for the simple group of 168 elements.

From sliding tiles to simple groups

The sliding tiles of the notorious "fifteen puzzle" (arranged in a 4×4 array with one square missing) are an object lesson in parity: Which permutations of the numbered tiles can be achieved? Precisely the even permutations. Put another way, the moves of the fifteen puzzle generate the alternating group A_{15}. Aaron Archer [1999] gives us a tidy proof of this folkloric fact.

R. M. Wilson [1974] considers tile-sliding games on arbitrary graphs as a generalization of the fifteen puzzle, and classifies the permutation groups which can be generated by these games. Briefly, the permutation group for a tile-sliding game on a graph with k vertices is generally either the alternating group A_{k-1} if the graph is bipartite or the full symmetric group S_{k-1} if it is not. There is only one interesting exception, a 7-vertex graph which generates a group of just 120 permutations. Wilson presents this group as $\mathrm{PGL}(2, \mathbb{F}_5)$, the group of Möbius transformations over the field of five elements, but it is isomorphic to the symmetric group S_5. Fink and Guy [2009] give a thorough discussion of this interesting, exceptional case, which they refer to as the "tricky six" puzzle.

John Conway [1997; 2006] uses a tile-sliding game on a 13-point projective plane to generate the Mathieu group M_{12}. This game is dubbed M_{13}. In M_{13}, moving a tile to the open point also requires swapping two other tiles at the same time. However, the rules of the game are specific to the projective plane on which it is played, and it does not generalize in any obvious way to a family of games on larger projective planes or other combinatorial structures.

MSC2010: 20B15, 91A43.
Keywords: simple group, permutation group, primitive group, graph theory.

Here, we consider a permutation game which can be played on any 3-regular graph, using a "slide-and-swap" rule inspired by Conway's M_{13}. We classify the resulting permutation groups and obtain a general result similar to Wilson's theorem, with just one interesting exceptional case. The exception gives a novel and elementary presentation for the simple group of 168 elements.

Like Rubik's cube and the other permutation games we have mentioned, ours can be treated purely as a puzzle, where you scramble pieces by moving them about and then try to return them to their initial configuration — or, perhaps, try to achieve some other "goal" configuration that has been posed as a problem. We have made a playable version of the game [Siehler 2011] as an aid to understanding the rules. In this article, we do not give any algorithms for unscrambling the pieces on a given graph, so solving the puzzle in that sense should remain an enjoyable challenge if you are so inclined.

Rules of the slide-and-swap game

Let Γ be a 3-regular graph (which we may as well assume to be connected) with labeled vertices. The game on Γ begins with one vertex uncovered and each of the remaining vertices covered by a tile with the same label. At any stage in the game, you make a move as follows: Choose a vertex v adjacent to the current uncovered vertex and slide the tile on v into the uncovered position (uncovering v in the process); and, at the same time, swap the tiles on the other two neighbors of v. See the first two diagrams in Figure 1, representing a single move where a tile slides from vertex b into the uncovered vertex and the tiles on the other two neighbors of vertex b are swapped.

We denote the move of sliding a tile from vertex b to vertex a by $[a, b]$. If this seems counterintuitive, think of the "hole" itself as a special blank tile which moves along the vertices in the order that they are listed when the move is played. Longer move sequences are expressed similarly: $[a, b, c, \ldots]$ is the move sequence which starts with the hole on vertex a, moves it to vertex b, from there to vertex c, and so on. Figure 1 shows a sequence of four moves played on the 8-vertex cube graph. The resulting permutation of tiles can be expressed in cycle form as $(b\ g\ c\ h\ d\ f\ e)$ — when this sequence is played, the tile which begins on vertex b moves to vertex g; the tile on vertex g moves to vertex c; and so on.

If $P = a, b, c, \ldots$ is a path in Γ, then "playing P" means playing the move sequence $[a, b, c, \ldots]$. In terms of tile movements, that means first moving the tile on vertex b to the vacant vertex a, then moving the tile on vertex c to the newly vacated vertex b, and so on, with the accompanying swaps at each step. The hole itself proceeds from a to b to c and so on, in the order they are listed.

Unlike the fifteen puzzle or M_{13}, the scrambling that happens as a result of

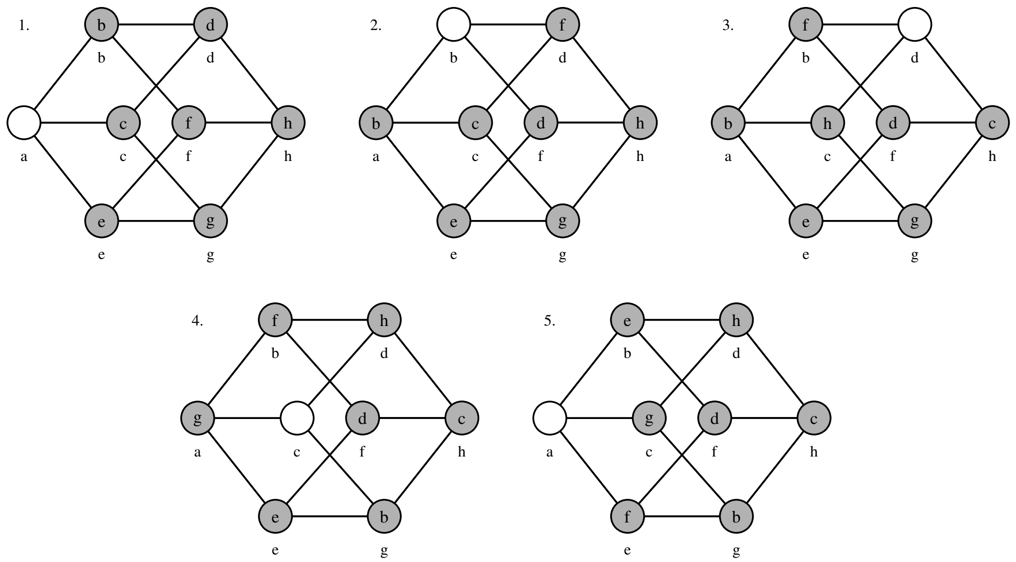

Figure 1. $[a, b, d, c, a] = (b\ g\ c\ h\ d\ f\ e)$ on the 8-vertex cube graph.

playing a path P in this game cannot be undone simply by playing its reverse (meaning the same path as P, just traversed in the opposite direction). However, we do have a basic result about invertibility:

Proposition 1. *Any legal move sequence in the slide-and-swap game can be undone (returning the tiles to their position before the sequence was played) by another legal move sequence.*

Proof. Suppose vertex x_0 is initially empty and we permute the tiles by playing $P = [x_0, x_1, \ldots, x_n]$. Let $R = [x_n, x_{n-1}, \ldots, x_0]$. The result of P followed by R is a permutation (of finite order, since there are only finitely many tiles!) which returns the hole to x_0. From that point we can play "P followed by R" repeatedly until all the tiles (and the hole) have returned to their initial position.

You may also note that a single slide $[x_0, x_1]$ can be undone by $[x_1, x_0, x_1, x_0]$. Consequently, a longer move sequence $[x_0, x_1, \ldots, x_n]$ can be undone one step at a time, as follows: First undo the final slide by playing $[x_n, x_{n-1}, x_n, x_{n-1}]$, then undo the one before that by $[x_{n-1}, x_{n-2}, x_{n-1}, x_{n-2}]$, and so on. □

Slide-and-swap loop groups

Now, suppose we begin a game on a connected, 3-regular graph Γ with k vertices, vertex a initially uncovered. The permutations which return the hole to its initial vertex (as in Figure 1) form a group under composition. We call this the *loop group* for Γ based at a and denote it $\mathcal{G}_a$. Since the basic moves in the game are all double transpositions, $\mathcal{G}_a$ is always a subgroup of the alternating group A_{k-1}. The notation suggests that the loop group depends on the choice of the initial uncovered vertex, but up to isomorphism the choice does not matter.

Proposition 2. *For any vertices a and b in* Γ, *the groups* $\mathcal{G}_a$ *and* $\mathcal{G}_b$ *are isomorphic.*

Proof. Since Γ is connected, choose a path P from a to b and let π be the permutation induced by playing P. The mapping $\alpha \mapsto \pi\alpha\pi^{-1}$ defines a homomorphism from $\mathcal{G}_a$ to $\mathcal{G}_b$. This homomorphism has $\beta \mapsto \pi^{-1}\beta\pi$ as its inverse, so the two groups are isomorphic (indeed, they are conjugate inside the symmetric group S_k). □

For this reason, we will henceforth omit any reference to the uncovered vertex and refer to *the* loop group associated to a graph.

Proposition 3. *The loop group of the tetrahedron* (*that is*, *the complete graph on* 4 *vertices*) *is trivial.*

The proof of this is left as an exercise. Larger graphs generate nontrivial groups, however, and a natural algebraic problem is to determine, up to isomorphism, which permutation groups can be realized as slide-and-swap loop groups. That problem is entirely resolved by the following theorems, which we will prove in the subsequent sections.

Notation. From now on, Γ will always denote a connected, 3-regular graph on k vertices, and $\mathcal{G}$ will denote its loop group (with the understanding that the choice of empty vertex doesn't matter).

Theorem 1. *If* Γ *is not the cube or tetrahedron*, *then* $\mathcal{G}$ *is isomorphic to the alternating group* A_{k-1}.

Theorem 2. *The loop group of the cube is isomorphic to* $\mathrm{GL}(3, \mathbb{F}_2)$, *the simple group of* 168 *elements.*

The resemblance to Wilson's results for "ordinary" tile-sliding games on graphs seems uncanny.

Fundamental terms and propositions

Dixon's problem book [1973] is a handy reference for the elementary theory of permutation groups, and the material is developed in depth in [Dixon and Mortimer 1996]. Here, we need only a few basic definitions and properties.

Let G be a group of permutations on a set X. The *orbit* of an element $x \in X$ is $\{\sigma(x) \mid \sigma \in G\}$. These orbits form a partition of X. If there is only one orbit (which contains all the elements of X), then G is said to be *transitive*.

A nonempty set $B \subseteq X$ is called a *block* for G if for every $\sigma \in G$, either $\sigma(B) = B$ or $\sigma(B) \cap B = \varnothing$. The set X itself is evidently a block, as are all singleton subsets of X; these are called *trivial* blocks. G is said to be *primitive* if all blocks for G are trivial; otherwise, if nontrivial blocks exist, G is *imprimitive*.

If G is transitive and B is any block, then the sets $\sigma(B)$, where $\sigma \in G$, partition X into disjoint, nonempty sets, each of which is a block. Such a partition of X is called a *system of imprimitivity* for G.

We consider $\mathscr{G}$ to be a group of permutations on the nonempty vertices of Γ. The following proposition is less obvious for slide-and-swap games than it is for ordinary tile-sliding games. Some time spent with a playable version of the game [Siehler 2011], trying to move a given tile to a given vertex, may be helpful in understanding the difficulties.

Proposition 4. *If* Γ *is not the tetrahedron*, *then* $\mathscr{G}$ *is transitive.*

Like Proposition 2, this is a useful fact to realize from the outset. It follows from the proof of Proposition 9, however, so we do not include a separate proof at this point. Our proof of Theorem 1 depends on the following general result:

Proposition 5 [Wilson 1974]. *Let G be a transitive permutation group on a set X and suppose that G contains a* 3*-cycle. If G is primitive*, *then G contains the alternating group on X.*

Generating the alternating group

The next few results show that, in general, our loop groups satisfy the hypotheses of Proposition 5. First, we establish the presence of 3-cycles.

Proposition 6. *If* Γ *is not the tetrahedron or cube*, *then* $\mathscr{G}$ *contains* 3*-cycles.*

Proof. First, note that the isomorphism in Proposition 2 is realized by conjugation within the symmetric group, which preserves cycle types. For this reason, if the group based at any vertex contains a 3-cycle, then this is true at every vertex, and so we can choose the empty vertex at our convenience. In the following, the labels (x_1, x_2, and so on) name the vertices of the graph. It doesn't matter which tiles are on which vertices, except that the vertex labeled x_0 is the initially open vertex. We consider three cases.

Case 1. If Γ contains a triangle $\{x_0, x_1, x_2\}$, then (since Γ is not a tetrahedron) two vertices in the triangle must have distinct neighbors. Suppose x_0 and x_1 have neighbors x_3 and x_4, like this:

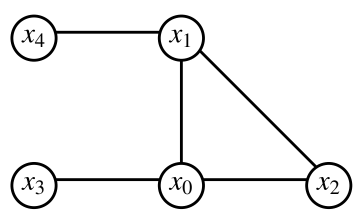

In this case, $[x_0, x_1, x_0] = (x_2\ x_4\ x_3)$.

Case 2. Suppose Γ does not contain a triangle, but contains a path x_0, x_1, x_2 where x_0 and x_2 have no neighbors in common except x_1:

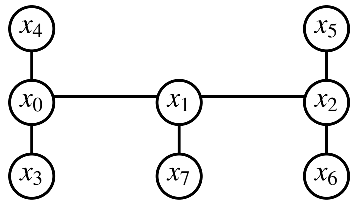

Then $[x_0, x_1, x_2, x_1, x_0]^2 = (x_1\ x_2\ x_7)$. Similarly, if Γ has no triangle but has a path x_0, x_1, x_2 where x_0 and x_2 have all three neighbors in common (say, $x_4 = x_5$ and x_3=x_6 in the figure), then $[x_0, x_1, x_2, x_1, x_0] = (x_1\ x_7\ x_2)$.

Case 3. The only remaining case to consider is a graph Γ with no triangles, in which the endpoints of every path of length two have exactly two neighbors in common. In *An atlas of graphs*, Read and Wilson [1998] show that there are only six 3-regular graphs of diameter less than three, including the tetrahedron, and none of them satisfy these hypotheses. These are all small graphs and the claim is easy to verify by inspection. Therefore we may assume that the diameter of Γ is at least 3. We claim that in this case Γ can only be the cube. Begin with a path a, b, c, d, where the distance from a to d is 3 (so there is no shorter path from a to d). Let v be the other common neighbor of a and c, and w the other common neighbor of b and d:

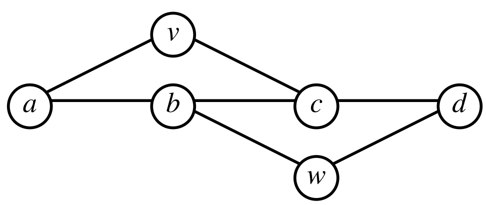

Now the path a, b, w implies that a and w have another common neighbor x. Similarly the path d, c, v implies that d and v must have another common neighbor y. This brings us to the following situation:

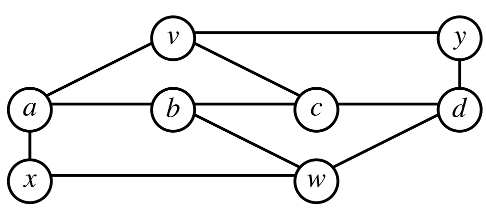

Since the graph is 3-regular, x needs another edge. If x were connected to some other vertex z not already shown, then the path z, x, a would imply that there is another vertex adjacent to both a and z. This would imply that either a, or one of the vertices adjacent to a in the preceding figure, has an additional edge, which is impossible; each of those vertices already has three edges. It follows that x and y are adjacent. At this point all the vertices have all three edges accounted for, and the resulting graph is the cube, as claimed. □

Proposition 7. *Suppose that for every pair of adjacent vertices x and y in Γ there is a p-cycle $\sigma \in \mathcal{G}$ where p is prime and $\sigma(x) = y$. Then $\mathcal{G}$ is primitive.*

Proof. Suppose the nonempty vertices are partitioned into blocks $B_1, B_2, \ldots, B_n$ forming a system of imprimitivity for $\mathcal{G}$. Let B_i and B_j be two blocks which contain adjacent vertices w_i and w_j, respectively. By hypothesis, we may choose a p-cycle $\sigma \in \mathcal{G}$ with $\sigma(w_i) = w_j$. Considered as a permutation of blocks, σ acts nontrivially (sending B_i to B_j), and since p is prime, this implies that σ acts with order p. However, σ only moves p vertices, so there can only be p blocks involved and only one vertex in each of those blocks, including B_i.

This argument can applied with any block playing the role of B_i and any adjacent block playing the role of B_j. Thus, since Γ is connected, either all blocks are singletons or there is only a single block. Since no nontrivial blocks are possible, $\mathcal{G}$ is primitive. □

Proposition 8. *Given any two adjacent, nonempty vertices x and y in Γ, there is a move sequence which leaves the tiles on x and y fixed while positioning the hole adjacent to one of those two vertices.*

Proof. If the empty vertex is already adjacent to x or y, no moves are needed. Otherwise, let Q be the shortest possible path beginning at the empty vertex, with the property that the ending vertex of Q has a distance of two from x or a distance of two from y. No point on Q can be adjacent to either x or y, since the previous point would have a distance of two, and a shorter path could be constructed. It follows that the permutation induced by playing Q leaves the tiles on x and y fixed. So, assume that we have played Q, ending with the hole on a vertex x_0 which is distance two from one of the given vertices — without loss of generality, say it's distance two from y. The goal now is to move the hole onto a vertex x_1 which is adjacent to y, still without disturbing the tiles on x and y. There are three local configurations to consider, and the proper moves to proceed in each case are shown in Figures 2–4.

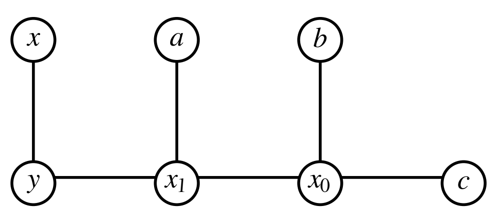

Figure 2. $[x_0, x_1, x_0, x_1] = (x_0\ x_1)(b\ c)$.

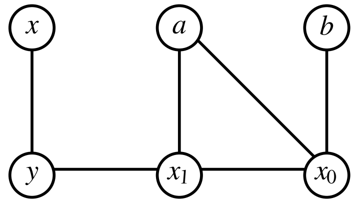

Figure 3. $[x_0, x_1, x_0, x_1, x_0, x_1] = (x_0\ x_1)(a\ b)$.

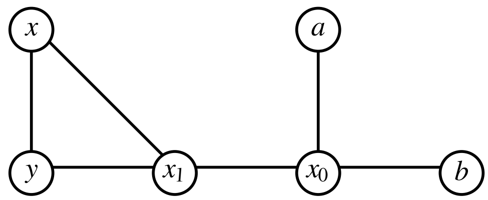

Figure 4. $[x_0, x_1, x_0, x_1] = (x_0\ x_1)(a\ b)$.

In each case, we get the desired result. Vertex x_1 (adjacent to y) is vacated, while the tiles on x and y remain fixed. We justify the claim that these three are the only cases as follows: Since x_0 is not adjacent to either x or y it must have two neighbors other than x, y or x_1. And x_1 must have one additional neighbor other than y or x_0. This additional neighbor may be x, or one of the neighbors of x_0, or a point distinct from both x and the neighbors of x_0, precisely the three cases we have considered. □

Proposition 9. *If* Γ *is not the tetrahedron, then the hypotheses of Proposition 7 are satisfied, and* $\mathscr{G}$ *is primitive.*

Proof. Once again, let x and y be any two adjacent, nonempty vertices in Γ. Assume the empty vertex is adjacent to y. Now, the possible configurations of the graph near x and y are summarized in the following figure, where x_0 is the empty vertex and x_1 is a vertex adjacent to y other than x and x_0.

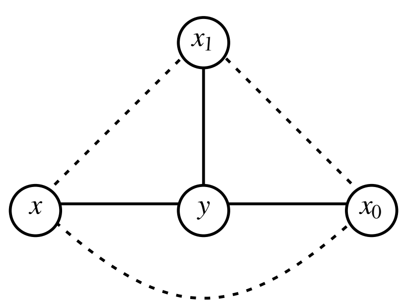

Each dotted edge may or may not exist. If all edges exist simultaneously then Γ is the tetrahedron. That leaves seven nontrivial cases to consider.

There are six cases in which at least one dotted edge exists; they are shown in

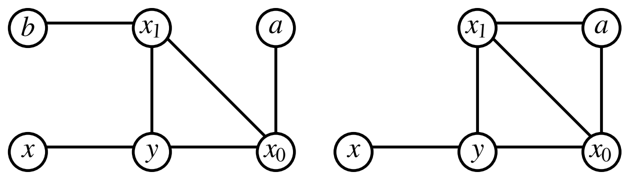

Figure 5. An x_0-x_1 edge. $[x_0, x_1, y, x_0]$ induces either $(x\ y\ b\ a\ x_1)$ or $(x\ y\ x_1)$.

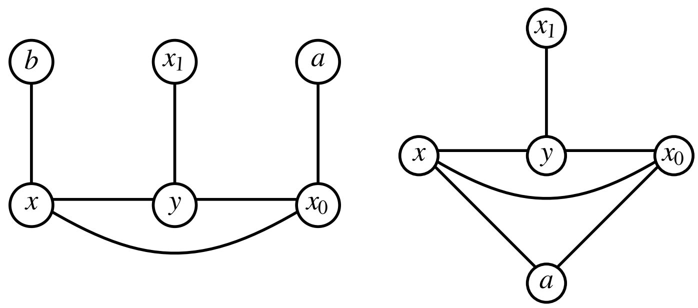

Figure 6. An x-x_0 edge. $[x_0, x, y, x_0]$ induces either $(x\ x_1\ y\ b\ a)$ or $(x\ x_1\ y)$.

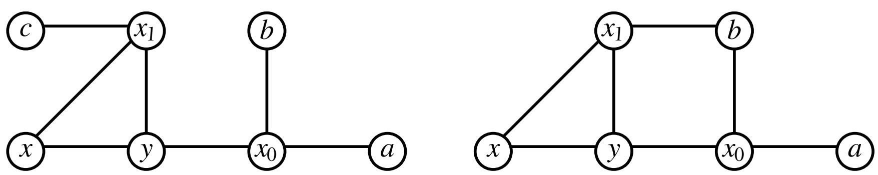

Figure 7. An x-x_1 edge. $[x_0, y, x_0, y, x_1, y, x_0]$ induces either $(x\ c\ y)$ or $(x\ a\ y)$.

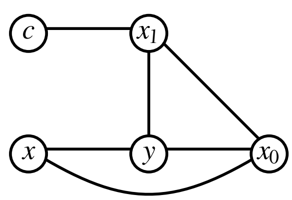

Figure 8. x-x_1 and x_0-x_1 edges. $[x_0, x_1, x_0] = (x\ y\ c)$.

Figures 5–10. In each case we exhibit a path which can be played to generate a cycle of prime length (either a 3- or a 5-cycle) sending x to y. Note that if a dotted edge is omitted, the vertices that edge connects must each have an edge to some other vertex not appearing in the figure at the bottom of previous page. These "extra" vertices are not necessarily distinct, so there are some subcases to be considered. In our figures, vertices not on the path or adjacent to a point on the path are not

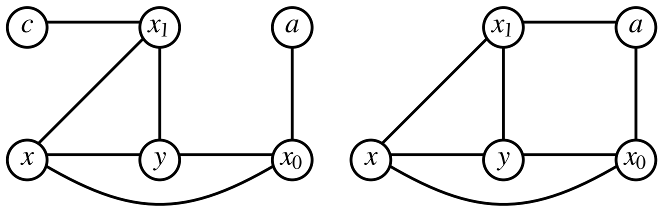

Figure 9. x-x_1 and x-x_0 edges. In the first case, $[x_0, y, x_1, y, x_0] = (x\ x_1\ c\ y\ a)$. In the second case, $[x_0, y, x_0, y, x_0, y, x_1, y, x_0] = (x\ y\ a)$.

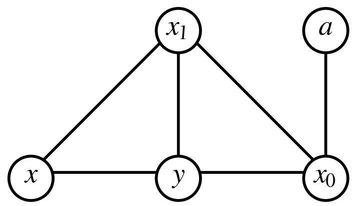

Figure 10. x-x_1 and x_0-x_1 edges. $[x_0, x_1, x_0] = (x\ a\ y)$.

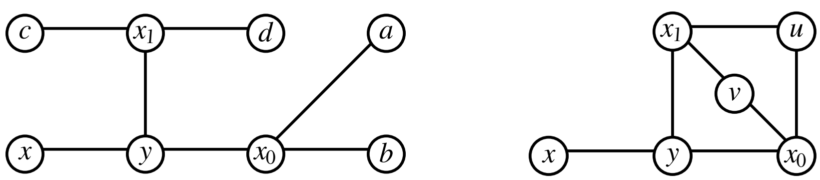

Figure 11. No edges among $\{x, x_0, x_1\}$. $[x_0, y, x_1, y, x_0]$ induces either $(a\ b)(c\ d)(x\ x_1\ y)$ (which can be squared to get the desired 3-cycle) or $(x\ x_1\ y)$.

drawn because they have no effect on the resulting permutation.

Figure 11 shows the case where none of the dotted edges are present, and x_1 and x_0 share 0 or 2 neighbors other than y.

That leaves the case where none of the dotted edges are present, and x_1 and x_0 have exactly one common neighbor u. The third neighbor of u is either one of the points on the graph other than y, or a point off the graph. So there are a total of four subcases to deal with in this case and they are shown in Figures 12–15.

In every case, we produce a cycle α of prime length which sends x to y. But we began with the provisional assumption that the initially empty vertex is adjacent to y. In general, however, Proposition 8 can be applied to produce a permutation σ which moves the empty vertex adjacent to y while leaving x and y fixed. Conjugating α by σ gives the desired cycle in $\mathcal{G}$. □

Remark. Since any vertex may be moved to any adjacent vertex by means of these cycles, $\mathcal{G}$ is transitive, as we asserted in Proposition 4.

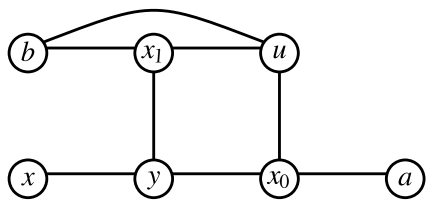

Figure 12. $[x_0, u, x_1, y, x_0] = (a\ u\ x\ y\ b)$.

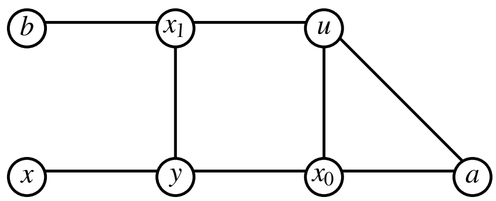

Figure 13. $[x_0, u, x_1, y, x_0] = (x\ y\ b\ x_1\ u)$.

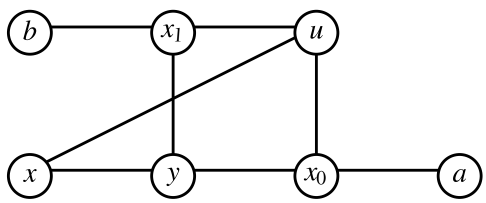

Figure 14. $[x_0, y, x_0, u, x_1, u, x_0] = (x\ u\ y\ b\ a)$.

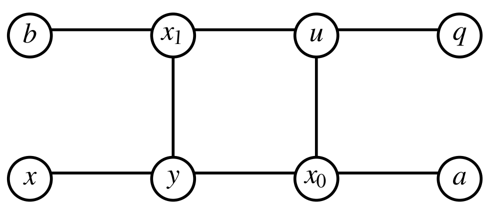

Figure 15. $[x_0, u, x_1, y, x_0] = (a\ u\ x\ y\ b\ x_1\ q)$.

Proof of Theorem 1. Our main result now follows quickly. If Γ is not the cube or tetrahedron, then Proposition 6 shows that $\mathscr{G}$ contains 3-cycles. Proposition 9 shows that $\mathscr{G}$ is transitive and primitive. Thus Proposition 5 applies to $\mathscr{G}$ and we conclude that $\mathscr{G}$ contains all even permutations of the nonempty vertices. □

The exceptional cube

Now, we analyze the loop group of the cube. Initially, a computer calculation revealed that this group has only 168 elements (instead of the expected $7!/2 = 2520$ in A_7). The number 168 is familiar to algebraists as the order of $\mathrm{GL}(3, \mathbb{F}_2)$, the group

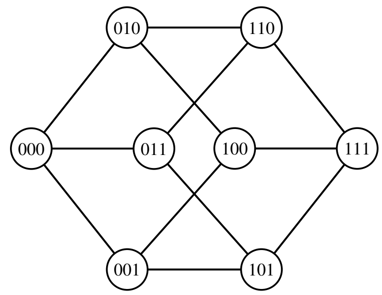

Figure 16. Vertices of the cube labeled with vectors of $\mathbb{F}_2^3$.

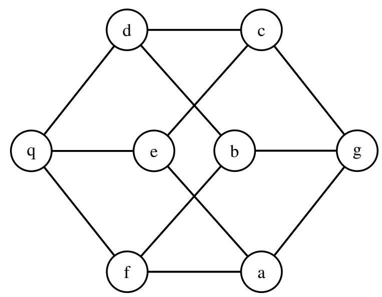

Figure 17. Vertices of the cube labeled with unknown vectors.

of invertible 3×3 matrices over the field of two elements, and the second-smallest nonabelian simple group.

To establish a connection between this group and the cube, we label the vertices with 3-dimensional vectors over $\mathbb{F}_2$, as in Figure 16. For brevity, we write a vector $\langle b_1, b_2, b_3 \rangle$ as a 3-bit binary string $b_1 b_2 b_3$. With such a labeling, moves in the game can be interpreted as permutations of $\mathbb{F}_2^3$. The particular labeling in Figure 16 has the property that the sum (mod 2, of course) of any tile together with its three adjacent tiles is zero, and we will call any arrangement of vector tiles with this property a *locally zero* arrangement.

Proposition 10. *The permutation of $\mathbb{F}_2^3$ induced by a single move on a cube with a locally zero arrangement of vectors is an affine transformation.*

Proof. Let $a, \ldots, g$ and q be the eight vectors of $\mathbb{F}_2^3$, labeling the vertices of the cube in a locally zero arrangement as in Figure 17. If we define

$$\alpha = a + q, \quad \beta = b + q, \quad \gamma = c + q,$$

the following additional relations follow quickly from the locally zero condition:

$$\beta + \gamma = d + q, \quad \alpha + \gamma = e + q, \quad \alpha + \beta = f + q, \quad \alpha + \beta + \gamma = g + q.$$

Note that all the linear combinations of α, β, and γ are distinct, so $\{\alpha, \beta, \gamma\}$ is a linearly independent set (and, in fact, a basis for $\mathbb{F}_2^3$).

Now, consider a slide-and-swap move. By symmetry we can assume that q is empty and we slide a tile into the hole from d. This induces a permutation φ with $\varphi(q) = d,\ \varphi(d) = q,\ \varphi(b) = c,\ \varphi(c) = b,$ and the other points remaining fixed. Define $\hat{\varphi}$ by $\hat{\varphi}(x) = \varphi(x + q) + d$. Applying $\hat{\varphi}$ to our basis elements gives

$$\hat{\varphi}(\alpha) = \alpha + \beta + \gamma, \quad \hat{\varphi}(\beta) = \beta, \quad \hat{\varphi}(\gamma) = \gamma,$$

from which we can verify the hard way (that is, by checking every linear combination of basis elements) that $\hat{\varphi}$ is linear:

$$\begin{aligned}
\hat{\varphi}(\alpha + \beta) &= \hat{\varphi}(f + q) = f + d = \alpha + \gamma = \hat{\varphi}(\alpha) + \hat{\varphi}(\beta),\\
\hat{\varphi}(\alpha + \gamma) &= \hat{\varphi}(e + q) = e + d = \alpha + \beta = \hat{\varphi}(\alpha) + \hat{\varphi}(\gamma),\\
\hat{\varphi}(\beta + \gamma) &= \hat{\varphi}(d + q) = q + d = \beta + \gamma = \hat{\varphi}(\beta) + \hat{\varphi}(\gamma),\\
\hat{\varphi}(\alpha + \beta + \gamma) &= \hat{\varphi}(g + q) = g + d = \qquad \alpha = \hat{\varphi}(\alpha) + \hat{\varphi}(\beta) + \hat{\varphi}(\gamma),
\end{aligned}$$

and of course $\hat{\varphi}(0) = \varphi(q) + d = d + d = 0$. Thus, with respect to the basis $\{\alpha, \beta, \gamma\}$, $\hat{\varphi}$ is represented by the matrix

$$M = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix},$$

and φ is described by the formula $\varphi(x) = M(x + q) + d$, or $\varphi(x) = Mx + (Mq + d)$, an affine transformation as claimed. □

Proof of Theorem 2. Begin the game with vertices labeled by vectors in a locally zero arrangement and the 000 vertex open. By the preceding proposition, any sequence of slides is a composition of affine transformations (which is again an affine transformation). The location of the hole always reveals the translation part of the transformation, so the loop group (corresponding to slides where the hole returns to 000) consists of linear transformations and is contained in $\mathrm{GL}(3, \mathbb{F}_2)$.

To complete the proof, we simply exhibit a few elements of the group. Returning to the cube in Figure 17 and supposing q is initially open, consider the elements

$$\begin{aligned}
[q, e, q, d, b, f, b, f, q] &= (d\ e)(a\ g\ b\ c),\\
[q, f, q] &= (d\ e)(a\ b),
\end{aligned}$$

which generate a dihedral group of eight elements. Also,

$$\begin{aligned}
[q, e, c, e, q] &= (f\ g\ e)(d\ c\ b),\\
[q, d, b, f, q] &= (f\ g\ b\ c\ e\ d\ a).
\end{aligned}$$

Subgroups of order 8, 3, and 7 imply a group of order at least 168, and so the loop group of the cube is not just contained in, but equal to $\mathrm{GL}(3, \mathbb{F}_2)$. □

Efficient solutions and other questions

In principle, the method of proof that we used to classify the loop groups could likely be turned into an algorithm for solving a scrambled puzzle on any given graph, since we show how to produce small, localized cycles at any point on the graph, which could be used to migrate pieces to their appropriate locations. In practice, this sort of solution takes many more moves than the optimal solution. In ordinary tile-sliding games, the problem of finding an optimal solution for a scrambled state is NP-complete [Goldreich 1984; Ratner and Warmuth 1990]. We do not know if the same is true for slide-and-swap games; this is a question for a future paper. The slide-and-swap variant is also unusual in that the number of moves required to achieve a position from start may be different from the number of moves required to return it to start. This aspect of the puzzle has no counterpart in other sliding or twisting permutation puzzles that we are familiar with, and the relationship between distance to and distance from start is worthy of some further analysis.

Conway, Elkies and Martin [Conway et al. 2006] have shown how to use duality of the projective plane to produce an outer automorphism of the Mathieu group M_{12}. It would be interesting if the symmetries of the cube and the slide-and-swap game rules allowed a similar construction of outer automorphisms for $\mathrm{GL}(3, \mathbb{F}_2)$, but we have not yet discovered how to do this, and it remains a subject for further investigation.

References

[Archer 1999] A. F. Archer, "A modern treatment of the 15 puzzle", *Amer. Math. Monthly* **106**:9 (1999), 793–799. MR 2001a:05004 Zbl 1007.00006

[Conway 1997] J. H. Conway, "M_{13}", pp. 1–11 in *Surveys in combinatorics, 1997*, edited by R. A. Bailey, London Math. Soc. Lecture Note Ser. **241**, Cambridge Univ. Press, 1997. MR 98i:20003 Zbl 0887.05002

[Conway et al. 2006] J. H. Conway, N. D. Elkies, and J. L. Martin, "The Mathieu group M_{12} and its pseudogroup extension M_{13}", *Experiment. Math.* **15**:2 (2006), 223–236. MR 2007k:20005 Zbl 1112.20003

[Dixon 1973] J. D. Dixon, *Problems in group theory*, Dover, 1973.

[Dixon and Mortimer 1996] J. D. Dixon and B. Mortimer, *Permutation groups*, Graduate Texts in Mathematics **163**, Springer, New York, 1996. MR 98m:20003 Zbl 0951.20001

[Fink and Guy 2009] A. Fink and R. Guy, "Rick's tricky six puzzle: S_5 sits specially in S_6", *Math. Mag.* **82**:2 (2009), 83–102. MR 2512593 Zbl 1223.05113

[Goldreich 1984] O. Goldreich, "Finding the shortest move-sequence in the graph-generalized 15-puzzle is NP-hard", 1984, Available at http://tinyurl/15puzz-pdf. preprint.

[Ratner and Warmuth 1990] D. Ratner and M. Warmuth, "The $(n^2 - 1)$-puzzle and related relocation problems", *J. Symbolic Comput.* **10**:2 (1990), 111–137. MR 91i:68138 Zbl 0704.68057