THE GROUP OF SYMMETRIES OF THE TOWER OF HANOI GRAPH

SOEUN PARK

ABSTRACT. The Tower of Hanoi problem with k pegs and n disks has been much studied via its associated graph, H_n^k , of legal states and moves. This has led naturally to structural questions about H_n^k itself. For example, what is its group of symmetries, $G(H_n^k)$? In this note, we prove for all $k \geq 3$ and $n \geq 1$ that $G(H_n^k) \cong S_k$, the group of permutations of k elements.

1. INTRODUCTION

The classical Tower of Hanoi problem consists of 3 wooden pegs and n disks with pairwise different diameters. The n disks are initially stacked on a single peg in order of decreasing size, from the largest at the bottom to the smallest at the top. (See Figure 1.) The goal is to move the tower of disks to another peg, moving one topmost disk at a time while never stacking a disk on a smaller one, by the minimal number of moves possible. The problem can be extended to the case with $k \geq 3$ pegs and n disks under the same set of rules.



FIGURE 1. Convention for labeling k pegs and n disks in the Tower of Hanoi game.

One can construct a graph modeling the Tower of Hanoi problem with k pegs and n disks, commonly called the Hanoi graph and denoted H_n^k . The Hanoi graph consists of vertices, or *legal states*, connected by

This work was undertaken at the Columbia University REU program supported by NSF grant DMS-0739392.



FIGURE 2. An example of a legal state.



FIGURE 3. A legal state that can be reached from Figure 2 by moving the topmost disk of peg 1 to peg 0.



FIGURE 4. An example of a perfect state.

edges between certain pairs of vertices. A legal state is any configuration of the n disks stacked on the k pegs so that no disk rests on a smaller disk. (See Figure 2.) A legal move connecting two legal states corresponds to taking the top disk off of any peg and moving it to another peg which contains no smaller disk. (See Figure 3.) Hence, if there exists a legal move between two legal states (or vertices), an edge is drawn between the two corresponding vertices in H_n^k . Moreover, a *perfect state* refers to any configuration where all n disks are stacked on a single peg. (See Figure 4.) Thus, we can restate the goal of the Tower of Hanoi game as passing from one perfect state to another through the shortest possible sequence of legal moves. The question about finding shortest paths has a well-known solution for k = 3 and the main open question in this subject is to find the length of the shortest paths for k > 3.

 H_n^3 has a particularly beautiful and simple recursive structure (see Figure 5), with fascinating connections to other seemingly-unrelated

mathematical objects, like Pascal's triangle [P94] or Sierpinski triangle. (See Figure 6.) It is therefore quite natural to ask whether various properties of H_n^3 hold for general H_n^k . Moreover, by learning about the structure of H_n^4 in particular, we hope to shed light on the question of shortest paths for k = 4.



FIGURE 5. The Hanoi graph with 3 pegs and 3 disks. Each 3-bit ternary string describes a legal state, as detailed in Definition 2.5.



FIGURE 6. The recursive pattern of the Sierpinski triangle.

For example, let $G(H_n^k)$ denote the group of automorphisms of H_n^k . It is straightforward to prove that $G(H_n^3) \cong S_3$, the group of permutations of 3 elements. This can be done by using the fact that geodesics (shortest paths with respect to the standard graph metric) between perfect states are unique in H_n^3 . This implies that any automorphism mapping the perfect states to themselves must also map each perfect state of every recursive substructure to itself. Therefore, by a simple inductive argument, any automorphism fixing perfect states must be the identity, restricting $G(H_n^k)$ to the group of permutations of the perfect states, i.e., S_3 .

This argument fails when k > 3, since it is a well known result that geodesics are no longer unique. The purpose of this note is to provide the following extension to $k \ge 3$.

Main Theorem. $G(H_n^k) \cong S_k$ for all $k \ge 3$ and all $n \ge 1$.

SOEUN PARK

Acknowledgements. I would like to thank Professors Jason Behrstock and Elisenda Grigsby for guiding me through completing my research topic during the Columbia University summer REU program by patiently reviewing my work and giving comments whenever needed. I also thank all the participants in the research group, Amy Cen, Michael Rand, Ian Spafford, and Andy Zhang as well as the graduate student advisor, Harold Sultan, for encouraging me to finish my work. Michael proved one of the lemmas necessary in this work and offered six of his well-drawn, unique figures of the Tower of Hanoi graphs, and Amy willingly let me use her drawing of the classical Tower of Hanoi graph. Andy and Ian inspired me with different approaches to the variations of the Tower of Hanoi problem that enabled me to see my project from different angles. Harold helped all of us by introducing LaTex and holding an extra research session every week. I would like to thank Columbia University math department, Professor Robert Friedman, and National Science Foundation for supporting the Columbia University 2008 summer REU program.

2. Definitions and Notations

A graph, Γ , consists of vertices and edges. We denote the set of vertices of Γ by $V(\Gamma)$. If there exists at least one edge whose boundary consists of a pair, $v_1, v_2 \in V(\Gamma)$, then v_1 and v_2 are said to be adjacent. We will also need the following standard definitions.

Definition 2.1. (Edge Matrix) Let Γ be a graph with $|V(\Gamma)| = N < \infty$, and a fixed ordering, $\beta = v_1, \ldots, v_N$, of $V(\Gamma)$. Then we can associate to Γ an $N \times N$ edge matrix, $[E(\Gamma)]_{\beta}$ (written with respect to the basis β), defined as follows. The entry e_{ij} in the *i*th row and *j*th column of $[E(\Gamma)]_{\beta}$ is the number of edges between v_i and v_j .

Definition 2.2. (Degree) The *degree* of a vertex, $v \in V(\Gamma)$, is the number of edges one of whose endpoints is v. Denote the degree of each vertex $v_i \in V(\Gamma)$ by $deg(v_i)$. In terms of the edge matrix $[E(\Gamma)]_{\beta}$, we have $deg(v_i) = \sum_{j=1}^{N} e_{ij}$.

Definition 2.3. (Automorphism) An *automorphism* of a graph, Γ , is an adjacency-preserving bijection $g: V(\Gamma) \to V(\Gamma)$. More precisely, if $\beta = v_1, \ldots, v_N$ is a chosen labeling of $V(\Gamma)$, and $[E(\Gamma)]_{\beta}$ is the corresponding edge matrix, then we denote by $[g(E(\Gamma))]_{g(\beta)}$ the edge matrix induced by g (written with respect to the labeling, $g(\beta)$, induced by g). Then g is a graph automorphism if and only if $[g(E(\Gamma))]_{g(\beta)} = [E(\Gamma)]_{\beta}$. It is clear that the set of automorphisms of any graph, Γ , form a group under composition. We denote by $G(\Gamma)$ the group of automorphisms of

4

5

the graph, Γ . Since an automorphism gives a symmetry of the graph, this group is sometimes called the *group of symmetries*.

Definition 2.4. A path, γ , between two vertices, $v, v' \in V(\Gamma)$ is a sequence of vertices v_0, \ldots, v_k satisfying the following properties.

- (1) $v_0 = v, v_k = v'$
- (2) v_j, v_{j+1} are connected by an edge for all $j = 0, \ldots, k-1$.

Then the *length* of such a path, γ , denoted by $l(\gamma)$, will be k. Define the distance d between v and v' by

$$d(v, v') = \min_{\gamma} \{ l(\gamma) \}$$

where γ ranges over all paths between v and v'. For every connected graph, Γ , d is well-defined, giving Γ the structure of a metric space. Moreover, one easily checks that any automorphism of Γ is an isometry of $V(\Gamma)$ with respect to this metric d.

Now we are ready to explicitly define the Tower of Hanoi graph, H_n^k , for k pegs and n disks.

Definition 2.5. (Legal States and Vertices) A vertex $v \in V(H_n^k)$ is an *n*-bit *k*-ary string, $a_{n-1}a_{n-2} \dots a_0$ with $a_i \in \{0, \dots, k-1\}$. Such a vertex corresponds to the legal state that has disk *i* on peg a_i , where disks are labeled from 0 to k-1, from smallest to largest. (See Figure 1.)

Remark 2.6. For a given *n*-bit *k*-ary string, all digits with the same a_i have to be stacked on peg a_i in a unique way, hence every *n*-bit *k*-ary string uniquely defines a legal state, and all legal states can be represented in this way. Hence, $|V(H_n^k)| = k^n$.

Definition 2.7. (Legal Moves) We denote by L_{ij} the legal move between peg *i* and *j*, from a vertex $a_{n-1}a_{n-2} \ldots a_0$. More precisely, consider the smallest disk of $a_{n-1}a_{n-2} \ldots a_0$ that is either on peg *i* or *j* (That is, the right-most occurrence of either *i* or *j* in $a_{n-1}a_{n-2} \ldots a_0$.) Call this term a_k . Then, $a_l \notin \{i, j\}$ for l < k. Let $\overline{a_k} = i$ if $a_k = j$, *j* if $a_k = i$. Then the legal move L_{ij} from the vertex $a_{n-1}a_{n-2} \ldots a_0$ between peg *i* and *j* is to the vertex given by $a_{n-1} \ldots a_{k+1} \overline{a_k} a_{k-1} \ldots a_0$. (I.e., Figure 2 and Figure 3 are connected by a legal move, L_{01} .)

Definition 2.8. (Edges and Edge Matrix) An *edge* is said to exist between vertices $v_1, v_2 \in V(H_n^k)$ if the corresponding legal states can be connected by a legal move. One easily finds that if a legal move exists between two vertices, it has to be unique. Therefore, every pair of adjacent vertices in H_n^k is connected by at most one edge. This implies that the entries of $[E(H_n^k)]_\beta$ are in $\{0, 1\}$.

SOEUN PARK

Remark 2.9. To prove that H_n^k is connected (and, hence, that $d(v_1, v_2)$) is well-defined for all $v_1, v_2 \in V(H_n^k)$, it suffices to show that there exists a path between a perfect state and an arbitrary state. It is wellknown that this is true for k = 3. We sketch the argument for general k > 3, assuming the result for k = 3. Without loss of generality, assume that we want to move a perfect state, 0_n , to an arbitrary state $v = a_{n-1}a_{n-2}\ldots a_0 \in V(H_n^k)$. First, by considering only the first 3 pegs (peg 0, 1, 2), we can distribute the tower of n disks on peg 0 to the pegs 0, 1, and 2 so that the disks on pegs 0 and 1 match those corresponding to the vertex v; and the disks on peg 2 are all the other ones. More precisely, 0_n has been transferred to an intermediate vertex $v_{0,1,2} = b_{n-1}b_{n-2}\dots b_0$ where $b_i = a_i$ if $a_i \in \{0,1\}$ and $b_i = 2$ otherwise. Repeat this process for pegs 1, 2, and 3 so that we can have disks on peg 1, 2 and 3 where the disks on peg 1 and 2 match those corresponding to the vertex v; and remaining disks (now other than the ones on peg (0, 1, 2) are all on peg 3. Therefore, 0_n has been transferred again from $v_{0,1,2}$ to another intermediate vertex $v_{0,1,2,3} = c_{n-1}c_{n-2}\ldots c_0$ where $c_i = a_i$ if $a_i \in \{0, 1, 2\}$ and $c_i = 3$ otherwise. By repeating this process a finite number of times, we obtain v in a finite number of legal moves.

Definition 2.10. (Corner vertices) A corner vertex of H_n^k is a vertex whose *n*-bit string corresponds to a perfect state. Denote the corner vertex corresponding to the *i*th peg by i_n .

Definition 2.11. (Substructures) A substructure, [i], is the set of vertices of H_n^k whose *n*-bit strings correspond to legal states whose largest disk lies on peg *i*. In particular, if $v \in [i]$, then $v = ia_{n-2} \dots a_0$.

Notation 2.12. Let S_k denote the group of permutations of the kelement set, $\{0, 1, \ldots, k-1\}$. For each $\sigma \in S_k$, define the map

$$g_{\sigma}: V(H_n^k) \to V(H_n^k)$$

by

$$g_{\sigma}(a_{n-1}\cdots a_1a_0) := \sigma(a_{n-1})\cdots \sigma(a_1)\sigma(a_0)$$

for $a_{n-1} \cdots a_1 a_0 \in V(H_n^k)$. Denote by $G(S_k)$ the set $\{g_\sigma \mid \sigma \in S_k\}$. As will be shown in the next section, $G(S_k)$ is actually a group and it is isomorphic to S_k .

3. Proof of Main Theorem

We prove the main theorem by showing that $G(S_k) \leq G(H_n^k)$ and $G(H_n^k) \leq G(S_k)$, which implies that $G(H_n^k) \cong G(S_k) \cong S_k$.

Step 1: $G(S_k) \leq G(H_n^k)$.

Proposition 3.1. For each $\sigma \in S_k$, g_{σ} is an automorphism of H_n^k . Hence, $G(S_k) \leq G(H_n^k)$.

Proof. For each $\sigma \in S_k$, we must show that g_{σ} is a bijection which preserves adjacency. To see that it is injective, we must show that $g_{\sigma}(a_{n-1} \ldots a_1 a_0) = g_{\sigma}(a_{n-1'} \ldots a_{1'} a_{0'})$ implies $a_{n-1} \ldots a_1 a_0 = a_{n-1'} \ldots a_{1'} a_{0'}$. For a given $\sigma \in S_k$, if $g_{\sigma}(a_{n-1} \ldots a_1 a_0) = g_{\sigma}(a_{n-1'} \ldots a_{1'} a_{0'})$, then $\sigma(a_{n-1}) \ldots \sigma(a_1) \sigma(a_0) = \sigma(a_{n-1'}) \ldots \sigma(a_{1'}) \sigma(a_{0'})$, which implies $\sigma(a_i) =$ $\sigma(a_{i'})$ for every $0 \le i \le n-1$. Moreover, since every $v = a_{n-1} \ldots a_1 a_0 \in$ $V(H_n^k)$ has a corresponding $v' = g_{\sigma^{-1}}(a_{n-1}) \ldots g_{\sigma^{-1}}(a_1) g_{\sigma^{-1}}(a_0) \in V(H_n^k)$ such that $g_{\sigma}(v') = v$, it is clear that g_{σ} is surjective. Hence, g_{σ} is bijective.

To prove that it preserves adjacency, we must show that $[E(H_n^k)]_{\beta} = [g_{\sigma}(E(H_n^k))]_{g_{\sigma}(\beta)}$ where β is the fixed basis of matrix $E(H_n^k)$ as in Definition 2.1. By Notation 2.8, $e_{ij} = 1$ implies that

$$v_i = a_{n-1} \dots a_{k+1} a_k a_{k-1} \dots a_0$$

$$v_j = a_{n-1} \dots a_{k+1} \overline{a_k} a_{k-1} \dots a_0$$

where $a_l \notin \{a_k, \overline{a_k}\}, \forall l < k$. Since g_{σ} is bijective, this is exactly equivalent to

$$g_{\sigma}(v_i) = \sigma(a_{n-1}) \dots \sigma(a_{k+1}) \sigma(a_k) \sigma(a_{k-1}) \dots \sigma(a_0)$$

$$g_{\sigma}(v_j) = \sigma(a_{n-1}) \dots \sigma(a_{k+1}) \sigma(\overline{a_k}) \sigma(a_{k-1}) \dots \sigma(a_0)$$

where $\sigma(a_l) \notin \{\sigma(a_k), \sigma(\overline{a_k})\}, \forall l < k$. Hence, $g_{\sigma}(e_{ij}) = 1$ iff $e_{ij} = 1$. Therefore, $\forall \sigma \in S_k, g_{\sigma}$ is an automorphism.

Step 2: $G(H_n^k) \leq G(S_k)$. We begin with a useful lemma.

Lemma 3.2. Every automorphism in $G(H_n^k)$ permutes the corner vertices, $\{0_n, 1_n, \ldots, (k-1)_n\}$. That is, for all $g \in G(H_n^k)$ and $i \in \{0, \ldots, k-1\}$, there exists a unique $j \in \{0, \ldots, k-1\}$ such that $g(i_n) = (j_n)$.

Proof. We use the fact that the degree of the corner vertices is strictly smaller than that of non-corner vertices. Since a graph automorphism preserves degree, it must therefore send corner vertices to corner vertices.

To see that the degree of each corner vertex is strictly smaller than the degree of each non-corner vertex, we compute the degree of a vertex in terms of the number of "topmost disks" of each vertex. For a given vertex, the topmost disk on each peg i is defined to be the smallest disk

SOEUN PARK

among those stacked on peg *i*. Therefore, the degree of a vertex is the sum of the number of legal moves each of its topmost disks can make. Since a corner vertex has only one topmost disk, which is disk 0, its degree is k - 1, the number of legal moves disk 0 can make. Now we must show that degree of each non-corner vertex is strictly larger than k - 1. If a vertex, v, is not a corner vertex, it has n disks distributed on at least 2 pegs. Hence, it has at least 2 topmost disks. Label the smallest and second smallest of those topmost disks, respectively b_0 and b_1 . Then b_0 can be moved to any other k - 1 pegs and b_1 to any other except where b_0 is stacked on, since b_0 is the only topmost disk that is smaller than b_1 . Therefore, every non-corner vertex has degree of at least (k-1) + (k-2) = 2k - 3, which is strictly larger than k - 1 for every $k \geq 3$.

To prove $G(H_n^k) \leq G(S_k)$ (and, hence, the main theorem), it suffices to prove the following proposition.

Proposition 3.3. If $g \in G(H_n^k)$ satisfies

$$g(i_n) = i_n \ \forall \ i \in \{0, \dots, k-1\},$$

then g is the identity automorphism. In other words, the only automorphism in $G(H_n^k)$ that fixes the corner vertices is the identity.

Proposition 3.3 implies $G(H_n^k) \leq G(S_k)$ because, by Lemma 3.2, every $g \in G(H_n^k)$ induces a permutation of the corner vertices, hence of the set $\{0, \ldots, k-1\}$. Therefore, there exists a $\sigma \in S_k$ that induces the same permutation of the set $\{0, \ldots, k-1\}$ as g does. Then, $g_{\sigma^{-1}} \circ g(0_n, 1_n, \ldots, (k-1)_n) = (0_n, 1_n, \ldots, (k-1)_n)$. Hence, as soon as Proposition 3.3 holds, $g_{\sigma^{-1}} \circ g$ has to be the identity automorphism, making $g = g_{\sigma} \in G(S_k)$. Hence $G(H_n^k) \leq G(S_k)$.

Proof of Proposition 3.3. We proceed by induction on n, fixing k.

Base case: n = 1. Trivial, since all vertices of H_1^k are corner vertices and the vertices are connected to one another by a single edge. Hence, any automorphism which fixes all the corner vertices fixes all vertices as well as edges, thereby inducing the identity automorphism.

Inductive Step: Assume that the proposition holds for n - 1. To prove that this implies that the proposition holds for n, we will need three lemmas, each helping prove the subsequent one. The argument for the first lemma was told to us by Michael Rand.

Lemma 3.4. [R08] Any shortest path between a corner vertex, i_n , and an arbitrary vertex, v, involves moving the largest disk zero times if $v \in [i]$; once otherwise.

Proof. It is clear that if i_n and v are in the same substructure, [i], the shortest path between them does not need to move the largest disk at all. To show that the largest disk moves exactly once when $v \notin [i]$, we prove that any path which moves the largest disk twice can be made strictly shorter by instead moving the largest disk once. Since any shortest path between i_n and $v \notin [i]$ that moves the largest disk more than twice necessarily contains a shortest path between an intermediate vertex and i_n that involves moving the largest disk exactly twice, the argument that it can't be moved twice therefore implies that it can't be moved *more than* twice as well.

We will prove by contradiction that no shortest path between i_n and v involves moving the largest disk twice. Without loss of generality, assume that there is a shortest path from 0_n to $v \notin [0]$ moving the largest disk exactly twice—first to peg 1 and then to peg 2. Then we can write the path as a sequence of steps:

- (1) Move the n-1 smallest disks off of peg 0 (leaving peg 1 clear at the end).
- (2) Move the largest disk from 0 to 1.
- (3) Some number of moves on the n-1 smallest disks (may be 0) which leave peg 2 empty and peg 1 containing only the largest disk.
- (4) Move the largest disk from 1 to 2.
- (5) Some number of moves (maybe 0) to get to the vertex v.

We claimed that this is the shortest path from 0_n to v, but we can create an even shorter path as follows.

- (1) Do the same moves as in (1), but with the roles of pegs 1 and 2 switched.
- (2) Move the largest disk from 0 to 2.
- (3) Do the same moves as in (3), but with the roles of pegs 1 and 2 switched.
- (4) Repeat step (5) above.

The second sequence of steps is one legal move shorter than the initial sequence, hence contradicts the assumption. Therefore, if $v \notin [0]$, the shortest path between 0_n to v involves moving the largest disk exactly once.

Lemma 3.5. For all $v \in [i]$, $d(v, i_n) < d(v, j_n) \forall j \neq i$.

Proof. To simplify notation we re-index the pegs; doing this, we assume for the rest of the argument that i = 0 and j = 1. By Lemma 3.4, the shortest path, γ , from $v \in [0]$ to 1_n involves moving the largest disk exactly once. Therefore we can further assume that γ can be split into



FIGURE 7. The sketch of proof for Lemma 3.5. $l(\gamma_2) = d(v_1, 1_n) = d(v_0, 0_n)$ and triangle inequality is applied to find $d(v, 0_n) < d(v, 1_n)$.

three parts: γ_1 , an intermediate legal move, and γ_2 . More precisely, γ_1 is the path from v to $v_0 = 0a_{n-2} \dots a_1a_0$, $a_i \notin \{0, 1\}$; an intermediate legal move is from $v_0 = 0a_{n-2} \dots a_1a_0$ to $v_1 = 1a_{n-2} \dots a_1a_0$; finally, γ_2 is from $v_1 = 1a_{n-2} \dots a_1a_0$ to 1_n . (See Figure 7.) We can easily observe that $l(\gamma_2) = d(v_1, 1_n) = d(v_0, 0_n)$ since the relationship between configurations v_1 and 1_n , and that of v_0 and 0_n are exactly symmetric. Hence,

$$\begin{aligned} d(v, 0_n) &\leq d(v, v_0) + d(v_0, 0_n) & \text{(by the triangle inequality)} \\ &< d(v, v_0) + 1 + d(v_0, 0_n) \\ &= d(v, v_0) + 1 + d(v_1, 1_n) \\ &= d(v, 1_n). \end{aligned}$$

Lemma 3.6. If $g \in G(H_n^k)$ satisfies $g(i_n) = i_n$, then $g(v) \in [i]$ for all $v \in [i]$. I.e., g([i]) = [i] as a set.

Proof. For concreteness we assume that 0_n is the vertex which is fixed by hypothesis. Now assume, aiming for a contradiction, that there exists $v \in [0]$ such that $g(v) \in [i], i \neq 0$. By Lemma 3.2, there exists a unique $j \neq 0$ such that $g(j_n) = i_n$. Thus, both g(v) and $g(j_n)$ are in [i]. Thus, since $g(j_n) = i_n$, by Lemma 3.6 we have

$$d(g(v), g(j_n)) < d(g(v), g(k_n))$$

 $\forall k \neq j$. Since $\forall g \in G(H_n^k)$ is an isometry, the previous statement implies $d(v, j_n) < d(v, k_n), \ \forall k \neq j$. Since $0 \neq j$, it follows that $d(v, j_n) < d(v, 0_n)$ with $v \in [0]$, which contradicts Lemma 3.5.

Now assume that Proposition 3.3 holds for $G(H_{n-1}^k)$ for all $k \geq 3$, and let $g \in G(H_n^k)$ satisfy $g(i_n) = i_n$ for all $i \in \{0, \ldots, k-1\}$. We have to show that this implies that g is the identity automorphism of H_n^k . By Lemma 3.6, $g(i_n) = i_n$ for all $i \in \{0, \ldots, k-1\}$ implies g([i]) = [i] for all $i \in \{0, \ldots, k-1\}$. Thus, for all $i, g|_{[i]}$ is an automorphism of

$$iV(H_{n-1}^k) := \{ia_{n-2}\cdots a_1a_0 \mid a_{n-2}\cdots a_1a_0 \in V(H_{n-1}^k)\}$$

Since the leading entry, i, is fixed, the automorphism g restricted to [i] is an automorphism which we denote $g_i : V(H_{n-1}^k) \to V(H_{n-1}^k)$. Moreover, g_i satisfies $g_i(i_{n-1}) = i_{n-1}$.

By Lemma 3.2, for all $j \neq i$, there exists a unique $l \neq i$ such that $g_i(j_{n-1}) = l_{n-1}$, hence $g(ij_{n-1}) = il_{n-1}$. We will now show that j = l, hence $g(ij_{n-1}) = ij_{n-1}$ for all j. We note the observation that ij_{n-1} is never adjacent to any vertex in [j], but on the other hand, has an adjacent vertex, lj_{n-1} , in any other substructure $[l], \forall l \neq j$. Since g is an automorphism it preserves adjacency, hence the above observation implies that $g(ij_{n-1})$ is still never adjacent to any vertex in g([j]) = [j], but has an adjacent vertex in any other substructure $g([l]) = [l], \forall l \neq j$. Hence $g(ij_{n-1}) = ij_{n-1}$ since otherwise it must be adjacent to a vertex in [j]. Thus, $g(ij_{n-1}) = ij_{n-1}$ for all j. This implies that for all i, $g_i \in G(H_{n-1}^k)$ satisfies the assumption of Proposition 3.3. Hence, by the inductive hypothesis, g_i is the identity for all i, implying that gfixes all the vertices of H_n^k . Thus, since there is at most one edge between any pair of vertices, any automorphism which fixes all the vertices must fix all the edges as well. Thus, we have shown q is the identity automorphism.

Step 1 and **Step 2** are now proved. Hence, $G(H_n^k) \cong G(S_k) \cong S_k$, as desired.

References

[P94] D.G. Poole, The Towers and Triangles of Professor Claus (or, Pascal Knows Hanoi). Math. Mag., 67:323-344, 1994.

[R08] M. Rand, Personal communication (2008).

E-mail address: sp2394@columbia.edu