Generating Linear Extensions of Posets by Transpositions

Frank Ruskey *
Department of Computer Science
University of Victoria
Victoria, B.C. V8W 2Y2, Canada

Abstract

This paper considers the problem of listing all linear extensions of a partial order so that successive extensions differ by the transposition of a single pair of elements. A necessary condition is given for the case when the partial order is a forest. A necessary and sufficient condition is given for the case where the partial order consists of disjoint chains. Some open problems are mentioned.

1 Introduction

Many combinatorial objects can be represented by permutations subject to various restrictions. The set of linear extensions of a poset can be viewed as a set of permutations of the elements of the poset. If the Hasse diagram of the poset consists of two disjoint chains, then the linear extension permutations correspond to combinations. If the poset consists of disjoint chains, then the linear extension permutations correspond to multiset permutations. The extensions of the poset that is the product of a 2-element chain with an n-element chain correspond to "ballot sequences" of length $2n$ and thus to binary trees with n nodes. More generally, the extensions of a poset that is an ideal in the product of two chains correspond to the standard Young tableaux having a fixed shape; see Figure 1 (also, Knuth [9]).

In this paper we are concerned with the question of generating all linear extensions of a poset so that successively generated permutations differ by a single transposition. In connection with sorting algorithms, transpositions have been called "exchanges" or "interchanges" (Knuth [9]). For computer scientists, linear extensions of posets are more well-known as "topological sortings". A number of algorithms exist for generating all the topological sortings of an arbitrary poset (or directed acyclic graph). See,

*Research supported by the Natural Sciences and Engineering Research Council of Canada under grant A3379.

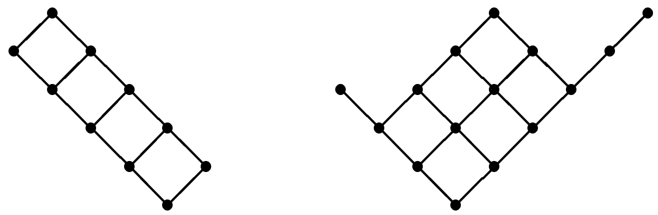

Figure 1: (a) Poset whose linear extensions correspond to binary trees. (b) Poset whose linear extensions correspond to Young tableau of shape 6,4,4,1.

for example, Kalvin and Varol [8], Varol and Rotem [24], and Knuth and Szwarcfiter [11]. Also, Knuth [10] considers the problem of generating the inverses of all permutations compatible with a forest partial order in lexicographic order. None of these papers address the problem of generating the linear extensions by transpositions.

A natural way to generate permutations is by transpositions. For unrestricted permutations the algorithm of Steinhaus [21], Johnson [7], and Trotter [23] is well-known and generates successive permutations by adjacent transpositions. In general, however, it is not always possible to generate all extensions by successive transpositions. For example, if the poset consists of two chains of lengths n and m, then the linear extensions can be generated by transpositions if and only if n and m are both odd (except for the trivial cases where $n = 1$ or $m = 1$), as shown independently by Buck and Wiedemann [2], Eades, Hickey, and Read [4], and Ruskey [17]. Here we generalize this result to permutations of a multiset.

There is a graph that naturally arises when generating extensions by transpositions. We call it the *transposition graph* of the poset. The vertices of this graph are the permutations that correspond to extensions of the poset and edges join those permutations that differ by a transposition. Note that this is always an induced subgraph of the transposition graph of all permutations. The problem of generating all the linear extensions of a poset is now transformed into that of finding a Hamiltonian path in the transposition graph of the poset. It has been observed by Ko and Ruskey [12] and Lehmer [13] in slightly different contexts that the transposition graph is bipartite. In the present context it is easy to see that the transposition graph is bipartite, since a transposition reverses the parity of a permutation. If the difference of the number of vertices in the two partite sets of the transposition graph is more than one, then it has no Hamiltonian path.

The transposition graph of a poset $\mathcal{P}$ is denoted $G(\mathcal{P})$. The *adjacent transposition graph* $G'(\mathcal{P})$ is the graph with the same vertex set as $G(\mathcal{P})$; however, edges only exist between permutations that differ by an *adjacent* transposition. These graphs are identical for the posets whose extensions correspond to combinations and to binary trees (in fact, for any poset of width two, since the presence of a nonadjacent transpositions requires a 3-element antichain).

An *inversion* of a permutation $e = e_1e_2\cdots e_n$ of $1, 2, \ldots, n$ is an ordered pair

(e_i, e_j) where $i < j$ and $e_i > e_j$. Note that an adjacent transposition changes the number of inversions by one. Let $inv(e)$ denote the number of inversions of e.

Suppose that $\mathcal{P}$ is a partial order on the set $\{1, 2, \ldots, n\}$ and that e and f are two of its linear extensions. We may assume without loss of generality that $f = 12 \cdots n$. We now claim that the distance in $G'(\mathcal{P})$ between e and f is $inv(e)$. Clearly it cannot be less. Furthermore, the elements of any adjacent inversion (e_i, e_{i+1}) are incomparable and thus can be transposed to obtain a linear extension that is closer to f.

By the preceding discussion $G'(\mathcal{P})$ is connected, but it implies more. Let $\mathcal{P}$ be a partial extension of $\mathcal{Q}$, meaning that every relation in $\mathcal{Q}$ occurs also in $\mathcal{P}$. Then $G'(\mathcal{P})$ is a convex subgraph of $G'(\mathcal{Q})$, where a subgraph H of a graph G is *convex* if, for every two vertices x and y of H, every vertex on a shortest path between x and y in G is also in H.

A poset is said to be *balanced* if it has the same number of even and odd permutations as extensions. A conjecture of Ruskey [18] states that the transposition graph of any balanced poset has a Hamiltonian path. The major result of this paper is that this conjecture is true for posets that consist of disjoint chains. This theorem is proved in section 3.

Theorem 1 *Multiset permutations can be generated by transpositions if and only if at least two types of elements have odd multiplicity, except for the trivial exceptions where there is one type of element or there are two types and one occurs with multiplicity one.*

Several parameters of the adjacent transposition graph $G'(\mathcal{P})$ are related to important parameters of the poset $\mathcal{P}$. The *jump number* of a linear extension $a_1 a_2 \cdots a_n$ is the number of indices $1 \leq i \leq n-1$ such that a_i and a_{i+1} are incomparable. Thus the jump number of a linear extension is the same as the degree of the corresponding vertex in $G'(\mathcal{P})$. The jump number of a poset is the minimum of the jump numbers of its linear extensions; so the jump number of a poset is equal to $\delta(G'(\mathcal{P}))$, where we use Bondy and Murty's notation $\delta(G)$ for the minimum vertex degree in G. Jump numbers are discussed in Chein and Habib [3], Rival [15], and Sysło [22].

The *bump number* of a linear extension $a_1 a_2 \cdots a_n$ is the number of indices $1 \leq i \leq n-1$ such that $a_i < a_{i+1}$. The bump number of a poset is the minimum of the bump numbers of its linear extensions. Thus $n-1$ minus the bump number of a poset is equal to the maximum vertex degree in $G'(\mathcal{P})$. Bump numbers are discussed in Fishburn and Gehrlein [5] and Zaguia [25].

In the next section, we discuss the problem for general posets. Here the problem is too unstructured and so not much is derived. In the third section we consider posets whose Hasse diagrams are forests. We derive a necessary condition for the adjacent transposition graph to have a Hamilton path; a condition that is conjectured to also be sufficient. In the fourth section we consider posets whose Hasse diagrams consist of disjoint chains. As noted above, the linear extensions of these posets correspond to multiset permutations. In Section 4.1 the necessary condition for the existence of a Hamilton path given by Ko and Ruskey [12] is shown to also be sufficient. This is the most significant result of the paper. In Section 4.2 we restrict our attention to adjacent transpositions. Partial results on the existence of Hamiltonian paths and cycles are stated. The final section mentions some open problems.

2 General Partial Orders

Let us review some definitions and introduce some further notation. A *partial order* is an antisymmetric, reflexive and transitive binary relation. A *poset*, $\mathcal{P}$, is a partial order $R(\mathcal{P})$ on a set $S(\mathcal{P})$. If $(a,b) \in R(\mathcal{P})$ then we write $a \leq b$; if also $a \neq b$ then we write $a < b$. For elements $a, b \in S(\mathcal{P})$, we say that b *covers* a and that (a,b) is a *cover relation* if $a < b$ and there is no $x \in S(\mathcal{P})$ such that $a < x < b$. An element of $S(\mathcal{P})$ is *minimal* if it covers no element; it is *maximal* if no element covers it. The set of all minimal elements is denoted $Min(\mathcal{P})$, and the set of all maximal elements is denoted $Max(\mathcal{P})$.

The *Hasse diagram* of a poset $\mathcal{P}$ is the diagram consisting of a graph whose vertices are the elements of $S(\mathcal{P})$ and whose edges are the cover relations, drawn in the plane so that if $a < b$ then b is above a (larger ordinate). Posets are often identified with their Hasse diagrams.

The number of elements of the poset $\mathcal{P}$ is denoted $|\mathcal{P}|$. Also, if x is in $S(\mathcal{P})$ then $\mathcal{P} - x$ denotes the poset $\mathcal{P}$ with the element x and all relations involving x removed.

A *linear extension* of a poset $\mathcal{P}$ is a permutation $a_1 a_2 \cdots a_n$ of the elements of $S(\mathcal{P})$ such that $a_i \leq a_j$ implies $i \leq j$. Let $e(\mathcal{P})$ denote the number of linear extensions of the poset $\mathcal{P}$. A recurrence relation for $e(\mathcal{P})$ is given below.

$$e(\mathcal{P}) = \begin{cases} 1 & \text{if } |\mathcal{P}| = 1 \\ \sum\limits_{x \in Min(\mathcal{P})} e(\mathcal{P} - x) & \text{otherwise} \end{cases} \tag{1}$$

Formula (1) is obtained by classifying the permutations according to their first element. As it stands, (1) is not an efficient means of calculating $e(\mathcal{P})$ since it essentially involves generating all linear extensions.

A formula similar to (1) for computing the difference between the number of vertices in the two partite sets of $G(\mathcal{P})$ can be derived. Define the *canonical permutation* to be $12 \cdots n$, where $n = |\mathcal{P}|$. We assume that the elements of the poset have been labeled in such a way that $12 \cdots n$ is a linear extension. Let the *parity difference* $d(\mathcal{P})$ be the number of even permutations that are extensions minus the number of odd permutations that are extensions. Stictly speaking, the parity difference is not solely a function of the poset since its sign depends on how the poset is labeled, but this will not cause us any problems. Below is a recurrence relation for $d(\mathcal{P})$.

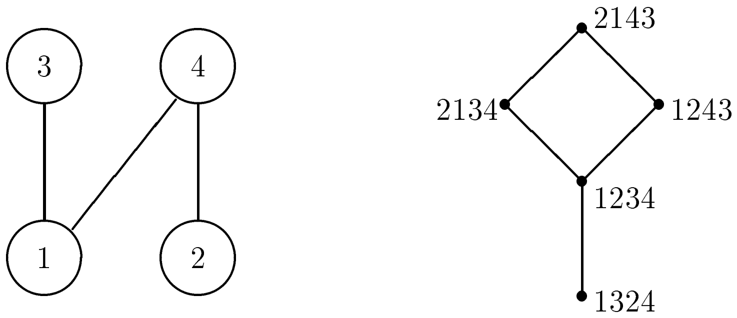

Figure 2: Example showing a poset and its adjacent transposition graph.

$$d(\mathcal{P}) = \begin{cases} 1 & \text{if } |\mathcal{P}| = 1 \\ \sum\limits_{x \in Min(\mathcal{P})} (-1)^{x-1} d(\mathcal{P} - x) & \text{otherwise} \end{cases} \tag{2}$$

In applying (2) it is assumed that the same ordering of the elements is used in defining the canonical permutation in $\mathcal{P} - x$ as was used in $\mathcal{P}$. Formula (2) is true since the parity difference for those permutations that begin with x is $d(\mathcal{P} - x)$ and it then takes $x - 1$ adjacent transpositions to bring the x into its proper place in the canonical permutation.

Recurrence relations (1) and (2) can be solved in special cases. Often this uses the fact that we could have used $Max(\mathcal{P})$ instead of $Min(\mathcal{P})$ and the recurrence relations would still be valid. This will be applied in the next section.

The example of Figure 2 illustrates some of the preceding discussion. The poset has been labeled in such a way that 1234 can be the canonical permutation. The transposition graph has 5 vertices (i.e., $e(\mathcal{P}) = 5$), the parity difference is -1 (i.e., $d(\mathcal{P}) = -1$) and the graph has a Hamiltonian path but no Hamiltonian cycle.

We now prove a small but useful lemma. The proof of this lemma serves to remind the reader of the "sweeping" idea behind the Steinhaus-Johnson-Trotter algorithm.

LEMMA 1 *Let z be incomparable to every element of $S(\mathcal{P})$. If $G(\mathcal{P})$ has a Hamiltonian path (cycle) then $G(\mathcal{P} \cup \{z\})$ has a Hamiltonian path (cycle).*

PROOF: Let $p_1, p_2, \ldots, p_N$ be the successive linear extensions along the Hamiltonian path (cycle) in $G(\mathcal{P})$. Suppose that $|\mathcal{P}| = n$. Then $p_1 z$ becomes $z p_1$ after n adjacent transpositions, where each transposition moves the z one position to the left. We say that the z has been "swept from right-to-left". Now change $z p_1$ to $z p_2$ and sweep the z from left-to-right through p_2. Continue in this manner: sweeping the z from right-left through p_i when i is odd, and from left-right when i is even. Since $G(\mathcal{P})$ is bipartite, if it has a Hamiltonian cycle then N is even. Thus the initial and final linear extensions of the Hamiltonian path of $G(\mathcal{P} \cup \{z\})$ are $p_1 z$ and $p_N z$; hence these two extensions differ by a transposition and there is a Hamiltonian cycle in $G(\mathcal{P} \cup \{z\})$. $\square$

Note that the same proof given above can be used to prove the corresponding lemma with G replaced by G'.

LEMMA 2 *If $G'(\mathcal{P})$ has a Hamiltonian path (cycle) then $G'(\mathcal{P} \cup \{z\})$ has a Hamiltonian path (cycle).*

Repeatedly applying the preceding lemma to the poset with n elements and no relations produces the Steinhaus-Johnson-Trotter algorithm for generating all $n!$ permutations of n elements.

3 Forests

If the Hasse diagram of a poset is a rooted forest then the recurrence relations (1) and (2) can be solved. We assume that the roots of the trees in the forest are the maximal elements of the poset and that the forest has been given a preorder labeling. A forest with $t+1$ trees is denoted

$$\mathbf{F} = T_0, T_1, \ldots, T_t$$

where each T_i is a tree. The forest that results when the root of a tree T is removed is denoted $s(T)$. Also, the following notation is introduced.

$$\mathbf{F}_i = T_0, T_1, \ldots, s(T_i), \ldots, T_t$$

As before, $e(\mathbf{F})$ denotes the number of linear extensions of the forest $\mathbf{F}$. A root is to the right of the elements in its subtrees in a linear extension. If the forest consists of chains (all vertices having one or zero children) then the scenario reduces to permutations of a multiset.

By summing over $Max(\mathbf{F})$ in formulas (1) and (2) we obtain (3) and (4) below.

$$e(\mathbf{F}) = \begin{cases} 1 & \text{if } k = 1, |T_1| = 1 \\ \sum_{i=1}^{k} e(\mathbf{F}_i) & \text{otherwise} \end{cases} \tag{3}$$

$$d(\mathbf{F}) = \begin{cases} 1 & \text{if } t = 0, |T_0| = 1 \\ \sum_{i=0}^{t} (-1)^{|T_0|+\cdots+|T_{i-1}|} d(\mathbf{F}_i) & \text{otherwise} \end{cases} \tag{4}$$

Suppose that there are n nodes in the forest and that the number of descendants of the ith node is d_i (a node is regarded as being a descendant of itself). It is an exercise in Knuth [9] (pg. 70) for the binary tree case to show that

$$e(\mathbf{F}) = \frac{n!}{d_1 d_2 \cdots d_n} \tag{5}$$

This equation also holds for rooted forests. It is similar to the "hook-length formula" for counting standard Young tableau [9].

We now show that there is a solution analogous to (5) to solve the parity difference recurrence relation (4).

DEFINITION 1 A forest $\mathbf{F}$ has a *pairing* if there is a 1-factor (perfect matching) in $\mathbf{F}$ or $\mathbf{F}_i$ for some i.

The pairing, if it exists, is unique and can be constructed "bottom-up", from the leaves towards the root(s).

DEFINITION 2 The *collapsed forest* $\overline{\mathbf{F}}$ of a forest $\mathbf{F}$ with a pairing is the forest obtained by contracting the matched edges of $\mathbf{F}$ and deleting the unmatched root, if there is one.

LEMMA 3 *For a forest* $\mathbf{F}$ *the parity difference is*

$$d(\mathbf{F}) = \begin{cases} e(\overline{\mathbf{F}}) & \textit{if } \mathbf{F} \textit{ has a pairing} \\ 0 & \textit{otherwise} \end{cases}$$

PROOF: We define an involution ϕ on the set of linear extensions of $\mathbf{F}$. Consider the elements of the permutation as being paired up from left to right. If $n = |\mathbf{F}|$ is even then the number of pairs is $n/2$; otherwise it is $(n-1)/2$. The involution takes the leftmost pair of a permutation π, say xy, for which y is not an ancestor of x (it cannot be a descendant), and returns the permutation which is identical to π except that the elements of the pair are transposed to yx. If there is no such pair then the involution returns π. Notice that if π is not a fixed point of ϕ then π and $\phi(\pi)$ have opposite parity.

Clearly, if ϕ has no fixed points then the parity difference is zero. Any fixed point induces the pairing of the forest $\mathbf{F}$ and every fixed point has the same parity. Thus the number of fixed points is equal to the number of linear extensions of the collapsed forest. If n is odd then the isolated root of the tree with an odd number of vertices must be the last element of any permutation that is a fixed point and so it can be omitted. $\square$

Some of the preceding ideas are illustrated in Figures 3 and 4. In Figure 3 all vertices except one can be paired, but the unpaired vertex is not a root vertex so the

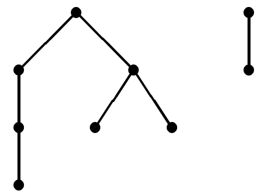

Figure 3: No pairing is possible; $d(\mathcal{P}) = 0$.

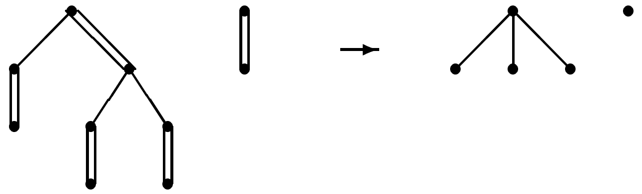

Figure 4: Pairing and collapsed forest; $d(\mathcal{P}) = 30$.

pairing is illegal. It is perhaps worth noting that the proof of Lemma 3 shows that $G'(\mathbf{F})$ always has a matching of cardinality equal to the size of the smaller partite set; if $d(\mathbf{F}) = 0$ then there is a perfect matching.

We believe that $G(\mathbf{F})$ has a Hamiltonian path if $d(\mathbf{F})$ is zero. If $\mathbf{F}$ is such that every vertex either is a leaf or has at least two children, then we can prove that $G(\mathbf{F})$ has a Hamiltonian path. The proof of this theorem gives us a warm-up for the arguments of the next section. This theorem may also be proved by using the results of Pruesse and Ruskey [14].

THEOREM 2 *If* $\mathbf{F}$ *is a forest poset of at least two trees in which no vertex has one child, then* $G(\mathbf{F})$ *has a Hamiltonian cycle.*

PROOF: The roots of the trees of the forest can be generated by adjacent transpositions by using the Steinhaus-Johnson-Trotter algorithm. Expand the forest top-down by adding children to an existing leaf vertex x. Let $z_1, z_2, \ldots, z_m$ with $m > 1$ be the children of x. Consider a generic permutation $\alpha x \beta$ before the children have been added. We wish to insert the m children in all possible ways into α. First insert the two children z_1 and z_2 such that the initial permutation is $z_1 z_2 \alpha x \beta$ and the final permutation is $z_2 z_1 \alpha x \beta$. That this is possible can be seen from Figures 5 and 6, where $n = |\alpha|$. In these figures the elements of α are represented by n y's, the vertices are all ways of inserting z_1 and z_2 into α. In the lower triangular portion z_1 is to the left of z_2 and in the upper triangular portion z_2 is to the left of z_1. Horizontal edges are transpositions of z_1 and a y, and vertical edges are transpositions of z_2 and a y. The diagonal edges are transpositions of z_1 and z_2.

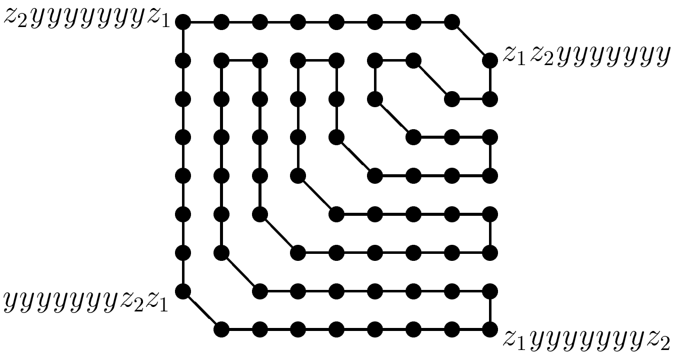

Figure 5: Hamiltonian cycle for n odd ($n = 7$).

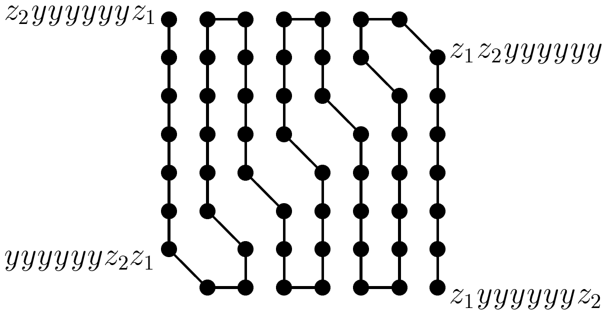

Figure 6: Hamiltonian cycle for n even ($n = 6$).

The remaining children $z_3, \ldots, z_m$ are added, a child at a time, using the Steinhaus-Johnson-Trotter "sweeping" idea; the initial permutation after adding the children is $z_m \cdots z_3 z_1 z_2 \alpha x \beta$ and the final permutation is $z_m \cdots z_3 z_2 z_1$. If $\alpha' x \beta'$ is the permutation that follows $\alpha x \beta$, then we proceed as before, except that the roles of z_1 and z_2 are reversed. Thus the initial permutation is $z_m \cdots z_3 z_2 z_1 \alpha' x \beta'$ and the final permutation is $z_m \cdots z_3 z_1 z_2 \alpha' x \beta'$. $\square$

Note that the proof of Theorem 2 does not provide a Hamiltonian cycle in $G'(\mathcal{P})$ because the transposition of $z_1 \alpha z_2$ to $z_2 \alpha z_1$ used in Figure 6 is not an adjacent transposition.

It would be interesting to know more about the structure of $G'(\mathbf{F})$. We can say something about the vertices of minimum degree. See also section 4.4.

Lemma 4 *If* $\mathbf{F}$ *is a forest with* l *leaves, then* $\delta(G'(\mathbf{F})) = l - 1$ *and there are* $l!$ *vertices with this minimum degree.*

Proof: As noted earlier, the degree of a vertex in $G'(\mathbf{F})$ equals the jump number of the corresponding linear extension. A linear extension $e_1 e_2 \cdots e_n$ is *greedy* if, for each $1 \leq i < n$, if there is a minimal element $x \in S(\mathcal{P}) - \{e_1, e_2, \ldots, e_i\}$ satisfying $x > e_i$, then $e_{i+1} > e_i$. For example, refering to Figure 2, the greedy extensions are 1324, 2134, and 2143. According to a result of Rival [15], there is a one-to-one correspondence between greedy linear extensions and linear extensions with minimum jump number, if the poset is N-free. Forests are obviously N-free. A forest with l leaves has $l!$ greedy linear extensions, and the jump number of the forest is $l - 1$. □

4 Multisets

In this section we consider posets that consist only of disjoint chains. A chain with k elements is denoted C_k. Since such a poset is completely determined by the number of elements in each of the chains, the notation $\mathbf{n} = n_0, n_1, \ldots, n_t$ is adopted to denote the poset consisting of $t + 1$ chains where the ith chain (counting from 0) has n_i elements.

Of course, $e(\mathbf{n})$ is the multinomial coefficient $n!/(n_0!n_1! \cdots n_t!)$, where $n = n_0 + n_1 + \cdots + n_t$. The recurrence relation for the parity difference for multisets was given and solved in [12] and [16], or can be derived from Lemma 3 of the previous section. Let $Odd(\mathbf{n})$ denote the number of n_i that are odd.

Lemma 5 *For multisets the parity difference is*

$$d(\mathbf{n}) = \begin{cases} e(\mathbf{m}) & \text{if } Odd(\mathbf{n}) \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

where $m_i = \lfloor n_i/2 \rfloor$ *and* $\mathbf{m} = m_0, m_1, \ldots, m_t$.

Lemma 5 establishes the necessity of the condition of Theorem 1; i.e., the parity difference will exceed one if $\mathbf{n}$ has fewer than 2 odd values (except for the trivial cases).

4.1 Hamiltonian Paths in G

In this subsection we show that the transposition graph $G(\mathbf{n})$ has a Hamiltonian path whenever $\mathbf{n}$ is balanced. The approach is inductive in t and the length of the chains. The base case is when $t = 2$, $n_0 = n_1 = 1$ (i.e., the poset consists of two isolated elements). We will show how to construct a Hamiltonian path in $G(\mathbf{n})$ that can be modified to obtain a Hamiltonian path in $G(\mathbf{n}')$, where $\mathbf{n}'$ is obtained from $\mathbf{n}$ by increasing the length of any chain by two or by adding a new chain of size at most two.

The actual Hamiltonian paths that are constructed might be called *sequentially adjacent*. This means that any transposition, say between elements x and y from the ith and jth chains, respectively, where $i < j$, satisfies the constraint that any element

z lying between x and y is from a chain of index greater than j. In particular, the elements from the last chain (of size n_t) only participate in adjacent transpositions.

It will prove useful to have some terminology to refer to the edges along some Hamiltonian paths in transposition graphs.

DEFINITION 3 Let $p_1, p_2, \cdots, p_N$ be a Hamiltonian path in the transposition graph of a balanced poset (thus N is even). The edges $[p_i, p_{i+1}]$ for i odd are called *expandable edges* or *e-edges*; the edges $[p_i, p_{i+1}]$ for i even are called *interface edges* or *i-edges*.

The permutations corresponding to vertices of the Hamiltonian path should be thought of as belonging to their respective expandable edges. The interface edges then only specify the transposition that is used in going from one e-edge to the next e-edge. In our subsequent constructions, the two permutations of an e-edge are replaced by larger sets of permutations, but the i-edges still specify the same transposition as before. In this way we think of interface edges as representing the interface between two sets of permutations.

An example will be used to explain our approach. Suppose that a Hamiltonian path in $G(5,3,2)$ has already been constructed and that we are now trying to construct a Hamiltonian path in $G(5,3,4)$. Let a, b, and z represent the elements from the chains of lengths 5, 3, and 2, respectively. As in [17], each linear extension p of $G(5,3,2)$ is "expanded" into the set of linear extensions of $G(5,3,4)$ that are identical to p when their two rightmost z's are deleted.

Suppose that $[p,q]$ is an e-edge, where $p = abzazabaab$ and $q = azbazabaab$. Inserting the two additional z's to the right of all existing z's in p in all possible ways produces the adjacent transposition graph shown on the left in Figure 7. The lower left corner is the permutation $abzazabaabzz$, the upper right vertex is $abzazzzabaab$, and the lower right corner is $abzazzabaabz$. In general, this graph is denoted I_n, where n is the number of elements to the right of the rightmost z. It is isomorphic to $G(2,n)$. For $n > 1$ there is no Hamiltonian path in I_n; this is the reason for considering pairs of vertices. Of course, the graph for inserting the two z's into q is also I_5. The adjacent transposition graph for inserting two z's into p and q is the Cartesian product graph of I_5 and an edge, which is denoted H_5 (in general, H_n). The vertices of the graph H_n will be drawn as shown in Figure 7. In order to keep the Hamiltonian paths shown in later figures uncluttered, no edges not on the path are drawn; they are left implicit. There are many Hamiltonian paths in H_n graphs. We are interested in paths that start in one of the copies of I_n and end in the other copy.

The only other type of e-edge that can occur is when the rightmost z is involved in the transposition, which conveniently did not happen above. For example, consider the edge $[p,q]$, where $p = abzazaabab$ and $q = abzzaaabab$. The adjacent transposition graph from inserting two additional z's into p and q is different from the graphs defined before. It is a I_n graph and a I_{n+1} graph with corresponding vertices joined by edges and is denoted J_n. The graph J_6 from our example is illustrated in Figure 8. As

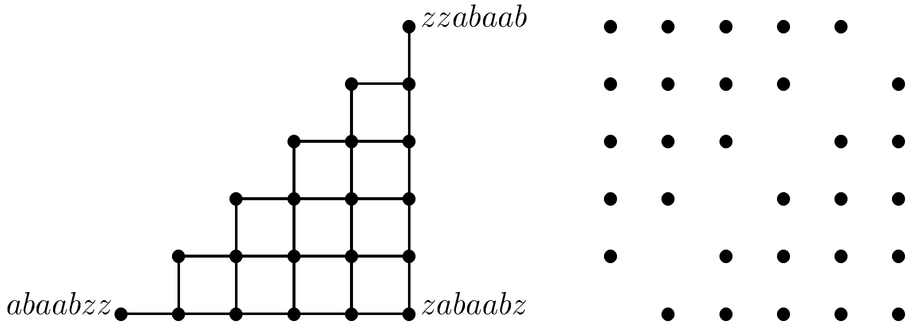

Figure 7: The graphs I_5 and H_4.

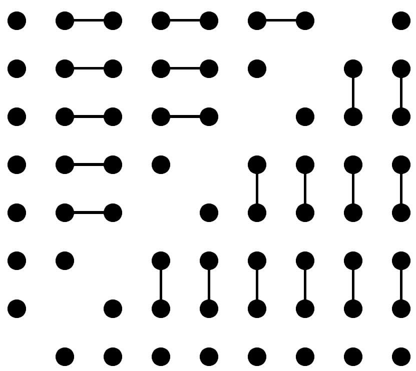

Figure 8: Forbidden e-edges in J_n (for $n = 6$).

before, the edges of the graph are usually not shown (ignore the edges shown for now). Note that there is a pendant (degree one) vertex in J_n and thus this vertex must be the initial or final vertex of any Hamiltonian path in J_n. We wish to find Hamiltonian paths that begin in the copy of I_{n+1} and end in the copy of I_n, or vice-versa.

Given Hamiltonian paths in H_n and J_n graphs, we would be able to produce a list of all permutations in $G(5,3,4)$, but the permutations at the interfaces of two successive expanded e-edges would not necessarily differ by a transposition. To ensure that they do differ by a transposition, we need to be careful about where the Hamiltonian paths in H_n and J_n graphs start and end. It will be necessary to define several different Hamiltonian paths.

The inductive hypothesis needs to be strengthened in order for the proof to proceed. Below we define a special type of Hamiltonian path in a transposition graph.

Definition 4 A Hamiltonian path $p_1, p_2, ..., p_N$ with N even in a transposition graph is a *z-proper path* if both of the following conditions are met:

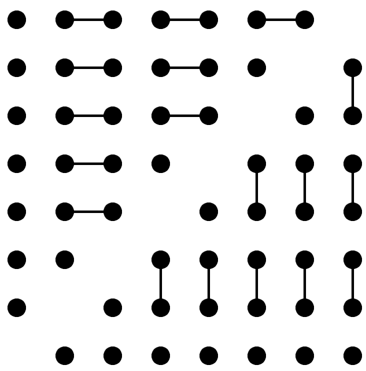

Figure 9: Forbidden e-edges in H_n (for $n = 6$).

1. **[The sequentially adjacent condition.]** If two successive permutations differ by a transposition involving z, then the transposition is an adjacent transposition.

2. **[The forbidden e-edge condition.]** There is no e-edge $[p, q]$ or $[q, p]$, where $p = \alpha xz\beta$, $q = \alpha zx\beta$, $|\beta|$ is odd and β contains no element greater than z.

Note that the parity of the number of vertices in J_n is the same as the parity of n and that the number of vertices in H_n is always even. In order for our approach to work the Hamiltonian paths constructed must start and end with e-edges and so the graphs must have an even number of vertices. Thus we can only use J_n graphs for n even. This is the reason for the "forbidden e-edge condition". Inductively, this places an additional constraint on the Hamiltonian paths. Movement of the rightmost z in an I_n graph (as drawn in Figure 7) corresponds to a vertical move. Thus certain vertical edges cannot be used (or horizontal edges in the reflected copies of I_n) as e-edges. The forbidden e-edges are shown in Figures 8 and 9. Those edges can (and will) be used as i-edges.

We adopt the convention that the lower case letters a, b, c, x, y, z represent single elements of the multiset permutations under consideration (all elements from the same chain are represented by the same letter). The letter z is always used to represent an element of the chain whose length is being increased or a new chain that is being added to the poset. The letters p and q always denote multiset permutations. Greek letters α, β, and γ represent sequences of zero or more elements of a permutation. The letter β is always a permutation that does not contain any z's.

Let us strengthen Lemma 1 in the context of proper paths. The corresponding result for adding a new chain of length two is implicit in our later results concerning the lengthening of chains by two.

Lemma 6 *If $G(\mathcal{P})$ has a Hamiltonian path (cycle) then $G(\mathcal{P} \cup C_1)$ also has a Hamiltonian path (cycle). If $e(\mathcal{P})$ is even then the resulting path (cycle) is proper with respect to the added element z.*

The different paths we use in H_n and J_n graphs are given in the two lemmas listed below. For any e-edge $[p, q]$ in which the transposition between p and q does not involve the final z, the transposition graph of all extensions arising by adding two z's is isomorphic to H_n, where n is the number of elements in p or q after the last z. Given p corresponding to the upper copy of I_n and q corresponding to the lower copy, the vertices are completely specified by indicating where the two new z's are inserted among the last n elements, and whether this insertion is in p (upper) or q (lower). Thus we indicate the vertices of H_n by $L:$ or $U:$ followed by a permutation of two z's and n x's. The transposition between p and q may trade two of the elements indicated by x's, or two elements not among the final n, or one of each type.

Lemma 7 *There are z-proper Hamiltonian paths in H_n from the initial vertex to the final vertex as indicated in the table below.*

type	conditions	initial	final
$H1$	n odd	$L{:}zzx^n$	$U{:}zzx^n$
$H2$	$n > 0$, even	$L{:}xzzx^{n-1}$	$U{:}xzzx^{n-1}$
$H3$	$n > 0$	$L{:}zzx^n$	$U{:}xzzx^{n-1}$
$H4$	$n > 0$	$L{:}xzzx^{n-1}$	$U{:}x^n zz$
$H5$	$n \geq 0$	$L{:}zzx^n$	$U{:}x^n zz$

Lemma 8 *For any even n, there are z-proper paths in J_n from the initial vertex to the final vertex as indicated in the table below.*

type	condition	initial	final
$J1$	n even	zzx^{n+1}	zzx^n
$J2$	$n > 0$, even	zzx^{n+1}	$xzzx^{n-1}$

Each of the preceding lemmas has a "dual" version where the path is followed in the opposite direction. In the dual version of Lemma 7 the path still goes from $L:$ representing p to $U:$ representing q and is indicated by interchanging both $L:$ and $U:$ and 'initial' with 'final'. We will denote the original version of the lemma by using a plus ($+$) superscript and the dual version by a minus ($-$) superscript. Path types $H1$ and $H2$ are self-dual and thus require no superscript. Here is the theorem that uses these lemmas.

Theorem 3 *If $G(n_0, n_1, \ldots, n_t)$ has a z-proper path, where z is the symbol used to represent the elements of the chain of length n_t, then $G(n_0, n_1, \ldots, n_t + 2)$ also has a z-proper path.*

PROOF: Let $p_1, p_2, \ldots, p_N$ be the permutations along a proper path in $G(n_0, n_1, \ldots, n_t)$. The two additional z's are inserted into all possible positions to the right of the rightmost z in p_i and p_{i+1} in every e-edge $[p_i, p_{i+1}]$ (for $i = 1, 3, 5, \ldots$).

Assume that all interfaces previous to the interface between p_{i+1} and p_{i+2} differ by a transposition (initially there is no interface to worry about). We need to consider the various types of interfaces that can occur.

The interface types are classified I, II, III, and IV, depending upon the placement and movement of the rightmost z within p_i and p_{i+1}. Class I consists of those e-edges where the rightmost z does not move and there are an even number of symbols to its right and Class II consists of those e-edges where the rightmost z does not move and there are an odd number of symbols to its right. Class III consists of those e-edges where the rightmost z moves one position left and Class IV consists of those e-edges where the rightmost z moves right. In both class III and class IV, the condition of starting with a z-proper path implies that the shorter number of symbols to the right of the rightmost z is even.

Each of the 18 possibilities is illustrated by an example from $G(5, 3, 2)$ in Table 1. The rightmost z is shown as a bold $\mathbf{z}$ to make it easier to see. Denote the initial permutation after expanding the e-edge e by $i(e)$ and the final permutation by $f(e)$. The cases where the rightmost z is the rightmost symbol of p_{i+1} or p_{i+2} are easily handled, and so we assume in the rest of the proof that there is at least one symbol to the right of the rightmost z in p_{i+1} and p_{i+2}.

$[p_i, p_{i+1}]$ **in Class I:**

If $e_1 = [p_i, p_{i+1}]$ is in Class I then it was expanded in the previous step by using path type $H2$, $H3$, $H4$, $H5$, or one of their duals. The final permutation after the expansion is then either $r = L\!:\!zzx^n$, $s = E\!:\!xzzx^{n-1}$, or $t = U\!:\!x^n zz$, where the E: indicates either L: or U:. We now show that we may choose $f(e_1) = s$ or $f(e_1) \neq s$ as needed for the interface with p_{i+2}.

If $f(e_1) \neq s$ then $H3^-$ or $H5^-$ was used to get $f(e_1) = r$, or $H4^+$ or $H5^+$ was used to get $f(e_1) = t$. If $H3^-$ was used then it is valid to use $H2$ instead since they both have the same initial permutation (i.e., $xzzx^{n-1}$). By using $H2$ we have $f(e_1) = s$. Similarly, if $H5^-$, $H4^+$, or $H5^+$ was used then use $H4^-$, $H2$, or $H3^+$ instead, respectively. By using the alternatives we always have $f(e_1) = s$.

If $f(e_1) = s$ then $H2$, $H3^+$, or $H4^-$ was used to expand e_1. Instead of $H2$ we can use $H4^+$ to get $f(e_1) = t$ (or $H3^-$ to get $f(e_1) = r$). Instead of $H3^+$ we can use $H5^+$ to get $f(e_1) = t$ and instead of $H4^-$ we can use $H5^+$ to get $f(e_1) = r$. By using the alternatives we always have $f(e_1) \neq s$. The argument of this paragraph is dual to the argument given in the previous paragraph.

Now consider the expansion of $e_2 = [p_{i+2}, p_{i+3}]$. For e_2 in cases I, II, III, or IV we assume that $f(e_1) = s$; for e_2 in case II$'$ we assume that $f(e_1) \neq s$. The following table indicates how e_2 can be expanded in such a way that $f(e_1)$ differs by a transposition from the initial permutation in the expansion of e_2.

		Class I	Class II	Class III	Class IV
	p_i	$aza\mathbf{z}babaab$	$azab\mathbf{z}abaab$	$azaba\mathbf{z}baab$	$az\mathbf{z}ababaab$
	p_{i+1}	$aza\mathbf{z}baabab$	$azab\mathbf{z}aabab$	$azab\mathbf{z}abaab$	$aza\mathbf{z}babaab$
I	p_{i+2}	$aza\mathbf{z}baaabb$	$aza\mathbf{z}baabab$	$aza\mathbf{z}babaab$	$aza\mathbf{z}baabab$
	p_{i+3}	$zaa\mathbf{z}baaabb$	$zaa\mathbf{z}baabab$	$zaa\mathbf{z}babaab$	$zaa\mathbf{z}baabab$
I$'$	p_{i+2}		$azaba\mathbf{z}abab$		
	p_{i+3}		$zaaba\mathbf{z}abab$		
II	p_{i+2}	$azab\mathbf{z}aabab$	$azab\mathbf{z}aaabb$	$azab\mathbf{z}aabab$	$azab\mathbf{z}abaab$
	p_{i+3}	$zaab\mathbf{z}aabab$	$zaab\mathbf{z}aaabb$	$zaab\mathbf{z}aabab$	$zaab\mathbf{z}abaab$
II$'$	p_{i+2}	$az\mathbf{z}abaabab$			
	p_{i+3}	$za\mathbf{z}abaabab$			
III	p_{i+2}	$aza\mathbf{z}baaabb$	$aza\mathbf{z}baabab$	$aza\mathbf{z}babaab$	$aza\mathbf{z}baabab$
	p_{i+3}	$az\mathbf{z}abaaabb$	$az\mathbf{z}abaabab$	$az\mathbf{z}ababaab$	$az\mathbf{z}abaabab$
IV	p_{i+2}	$azab\mathbf{z}aabab$	$azab\mathbf{z}aaabb$	$azab\mathbf{z}aabab$	$azab\mathbf{z}abaab$
	p_{i+3}	$azaba\mathbf{z}abab$	$azaba\mathbf{z}aabb$	$azaba\mathbf{z}abab$	$azaba\mathbf{z}baab$

Table 1: The 18 possibilities.

Case of e_2	I	II	II$'$	II$'$	III	IV
$f(e_1)$	s	s	r	t	s	s
To expand e_2	$H2$	$H1$	$H3^-$	$H4^-$	$J1^-$	$J1^+$
or	$H3^-$	$H3^+$	$H4^+$	$H5^-$		$J2^+$
or	$H4^+$	$H5^+$				

For the example of Class I to Class I given in Table 1, we could use the $H2$ path to obtain $i(e_2) = aza\mathbf{z}bzzaaabb$ and $f(e_2) = zaa\mathbf{z}bzzaaabb$, or the $H3^-$ path to obtain $f(e_2) = zaa\mathbf{z}zzbaaabb$, or the $H4^+$ path to obtain $f(e_2) = zaa\mathbf{z}baaabbzz$.

$[p_i, p_{i+1}]$ in Class II:

If $e_1 = [p_i, p_{i+1}]$ is in Class II then it was expanded in the previous step by using path type $H1$, $H3$, $H4$, or $H5$. The final permutation after the expansion is then either $r = E : zzx^n$, $s = E : xzzx^{n-1}$, or $t = U : x^n zz$. We can show, as was done above for e_1 in Class I, that we may choose $f(e_1) = r$ or $f(e_1) \neq r$ as needed for the interface with p_{i+2}.

Now consider the expansion of $e_2 = [p_{i+2}, p_{i+3}]$. For e_2 in cases I, II, III, or IV we assume that $f(e_1) = r$; for e_2 in case I$'$ we assume that $f(e_1) \neq r$.

Case of e_2	I	I$'$	I$'$	II	III	IV
$f(e_1)$	r	s	t	r	r	r
To expand e_2	$H2$	$H3^+$	$H4^-$	$H1$	$J2^-$	$J1^+$
or	$H3^-$	$H5^+$	$H5^-$	$H3^+$		$J2^+$
or	$H4^+$			$H5^+$		

$[p_i, p_{i+1}]$ **in Class III or IV:**

In Class III we must have $f(e_1) = zzx^{n+1}$. In Class IV we may assume, without loss of generality, that $f(e_1) = xzzx^{n-1}$ (i.e., that $J2^+$ was used in expanding e_1), n being even in either case. Note from the table that under this assumption the resulting $i(e_2)$ after the interface transposition following $f(e_1)$ does not depend on whether e_1 is in Class III or Class IV. For e_2 in Class I,II,III,IV, we have $i(e_2) = xzzx^{n-1}$, zzx^{n+1}, zzx^{n+1}, or $xzzx^{n-1}$, respectively, where n is even.

Case of e_2	I	II	III	IV
To expand e_2	$H2, H3^-, H4^+$	$H1, H3^+, H5^+$	$J2^+$	$J1^-, J2^-$

This completes the proof of Theorem 3. □

Theorem 1 is an immediate corollary of Theorem 3. We could have proven Theorem 3 by modifying the proof for $t = 1$ given in Ruskey [17]. The advantage of the proof given here is that it relies only on the statements of Lemmas 7 and 8, and not on the specific paths that were used to prove those lemmas; this allows us to present a complete proof of reasonable length. The proof in [17] relies on the specific paths that are used and many details are left to the reader to verify.

As in [17] our proof of Theorem 3 leads directly to an algorithm that generates all multiset permutations in constant average time. By *constant average time* we mean that the total number of computational steps divided by the number of multiset permutations generated is bounded by a constant.

We now indicate how to prove Lemmas 7 and 8. Once appropriate paths are exhibited in H_1, H_2, and J_2, the paths for n can be expanded into corresponding paths for $n+2$ to complete an iductive proof. Each proof is suggested by a figure illustrating the appropriate path in H_n or J_n for a specific value of n as shown in Figures 10, 11, 12, 13, and 14. Some paths depend on the parity of n. In the figures the e-edges are denoted by either thick edges or wavy edges. The thick edges correspond to transpositions not involving the rightmost z and the wavy edges correspond to transpositions involving the rightmost z. Note that the wavy edges avoid the forbidden e-edges of Figures 8 and 9. In H_0 we need only use paths of type $H5$, and in J_0 we need on use paths of type $J1$.

Although we have shown the existence of a Hamiltonian path in G, it remains an open question as to whether there is a Hamiltonian path in the spanning subgraph G'. We conjecture that there is a Hamiltonian path in this case also.

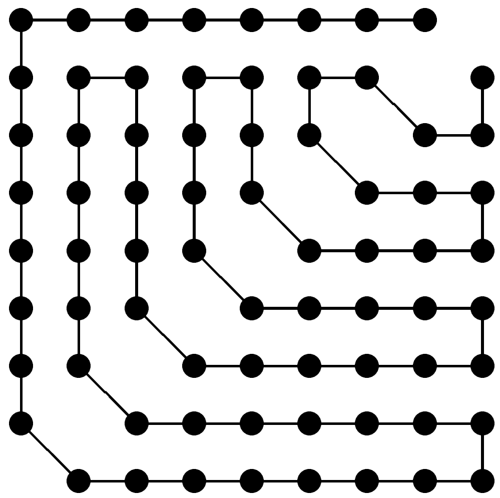

Figure 10: Hamiltonian path of type $H1$ ($n = 7$).

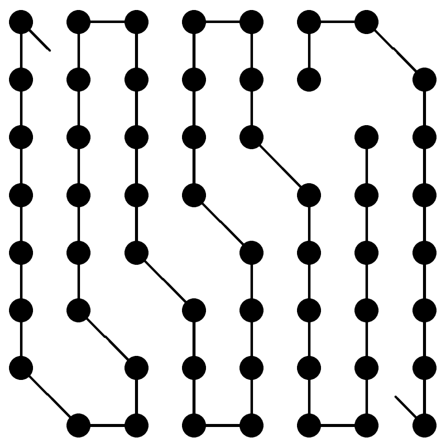

Figure 11: Hamiltonian path of type $H2$ ($n = 6$).

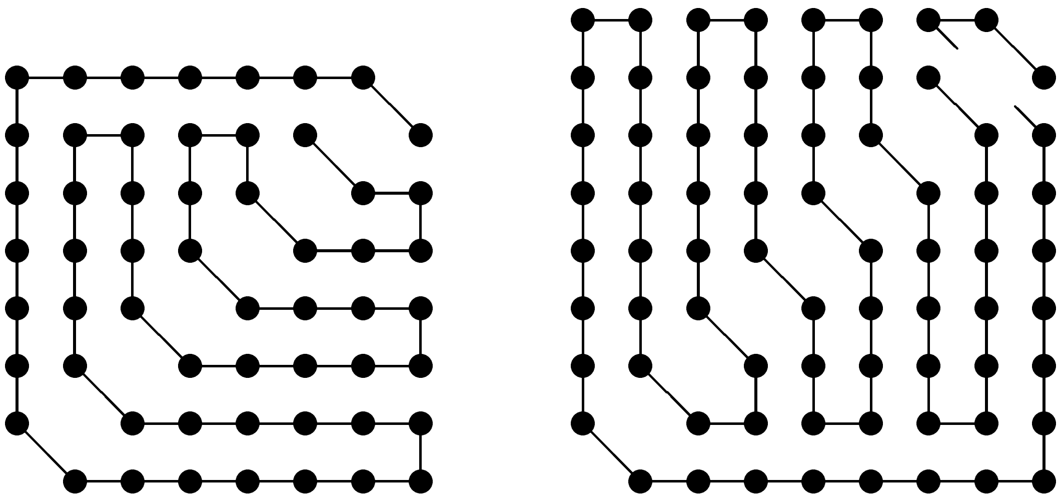

Figure 12: Hamiltonian path of type $H3$ ($n = 6, 7$).

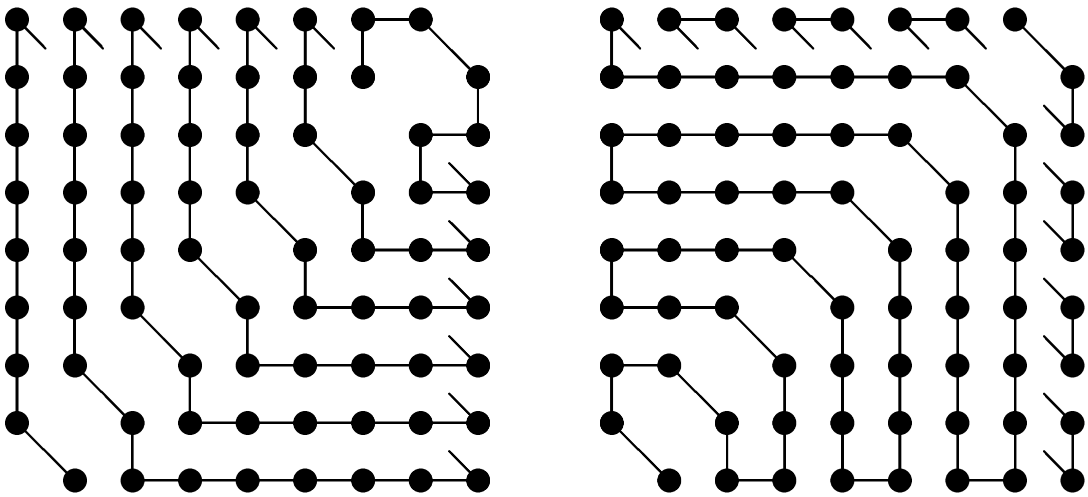

Figure 13: Hamiltonian paths of type $H4$ and $H5$ ($n = 7$).

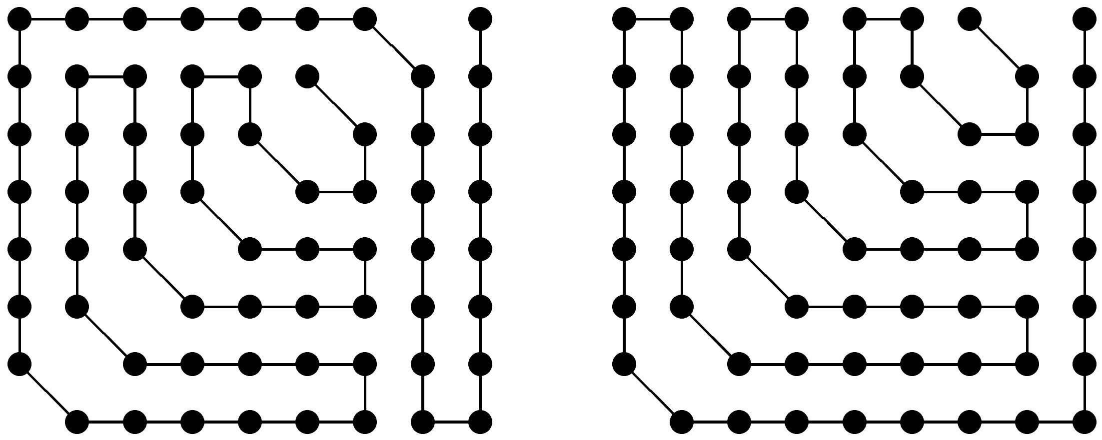

Figure 14: Hamiltonian path of type $J1$ and $J2$ ($n = 6$).

4.2 Hamiltonian Paths and Cycles in $G'(\mathbf{n})$

In this section, we provide partial results about Hamiltonian paths and cycles in $G'(\mathbf{n})$. We believe that $G(\mathbf{n})$ has a Hamiltonian cycle whenever $\mathbf{n}$ is balanced and there are at least three symbols (i.e., $t > 1$). As observed before, there cannot be a Hamiltonian cycle if $t = 1$, since then $G'(\mathbf{n})$ has two pendant vertices. Our first result shows that even when there are no pendant vertices, the adjacent transposition graph G' can fail to have a Hamiltonian cycle (although G has a Hamiltonian cycle for this example).

Lemma 9 *The graph $G'(1, 1, n)$ has a Hamiltonian cycle if and only if n is odd.*

Proof: The graph $G'(1, 1, n)$ consists of two copies of I_n. All vertices in one copy have the elements from the two chains of length one, call them x and y, in the same position relative to each other. The only other edges of $G'(1, 1, n)$ are between vertices of the form $\alpha xy\beta$ and $\alpha yx\beta$; call these *diagonal vertices*. See Figure 15. If n is odd then Figure 5 shows how to construct the Hamiltonian cycle.

If n is even, then the following argument shows that no Hamiltonian cycle exists. Note that there are $2(n + 1)$ diagonal vertices. If the $n(n + 2)/2$ vertices that are an odd distance from the nearest diagonal vertex are removed, then $n(n + 2)/2 + 1$ components remain. It is well known (see, e.g., Bondy and Murty [1], Theorem 4.2) that if more components remain than vertices that are removed, then the graph is non-Hamiltonian. Hence $G'(1, 1, n)$ with n even is not Hamiltonian. □

The graphs $G'(1, 1, n)$ with n even and the graphs $G(n, m) = G'(n, m)$ with n and m both odd are the only examples we know where a Hamiltonian path in $G'(\mathbf{n})$ exists

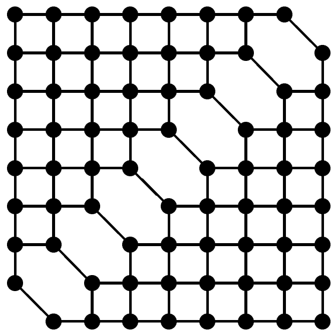

Figure 15: The graph $G'(1,1,n)$ (for $n = 7$).

but not a Hamiltonian cycle. We believe these graphs to be the only exceptions. In the remainder of this section we show that certain other adjacent transposition graphs have Hamiltonian cycles or paths.

LEMMA 10 *Let $\mathcal{P}$ be a balanced poset. If $G'(\mathcal{P})$ has a Hamiltonian path then $G'(\mathcal{P} \cup C_2)$ has a Hamiltonian path. If $|\mathcal{P}| \geq 3$ and $G'(\mathcal{P})$ has a Hamiltonian cycle then $G'(\mathcal{P} \cup C_2)$ has Hamiltonian cycle.*

PROOF: Assume that $\mathcal{P}$ is a poset that satisfies the conditions of the lemma. Pair up the succeeding permutations of the Hamiltonian path starting at some end to form e-edges. Let $[A, B]$ be an e-edge where $A = \alpha xy\beta$ and $B = \alpha yx\beta$. Inserting the C_2 elements, call them z, produces an adjacent transposition graph that is denoted $H_{n,k}$, where n is $|A|$ and k is $|\alpha x|$. The graph $H_{n,k}$ is the subgraph of H_n defined by deleting those "crossing" edges that join permutations which differ by a non-adjacent transposition of x and y; the vertices corresponding to these pairs are indicated by open circles in Figure 16 for $H_{6,2}$, and in Figure 17.

Let us use the same convention for refering to the vertices of $H_{n,k}$ as was used for H_n. We now show that the graph $H_{n,k}$ has a Hamiltonian cycle if $n \geq 3$. Furthermore, the cycle includes the edges $e_1 = [U : x^n zz, U : x^{n-1}zzx]$, $e_2 = [U : zzx^n, U : xzzx^{n-1}]$, and $e_3 = [L : x^n zz, U : x^n zz]$. The argument is inductive, but is presented informally in the figures described below. Edges e_1 and e_2 are used in the inductive construction and e_3 is used for the interfaces. Figure 17(a) shows that $H_{3,1} \approx H_{3,2}$ has a Hamiltonian cycle. The basic idea is that the cycle constructed in $H_{n,k}$ can be modified to obtain cycles in $H_{n+1,k}$ and $H_{n+1,k+1}$ as shown in Figure 18.

If $|\mathcal{P}| = 2$ then $\mathcal{P}$ consists of two unrelated elements and the Hamiltonian path is the one indicated in Figure 17(b). If $|\mathcal{P}| \geq 3$ then use the Hamiltonian cycles constructed above for each successive pair of vertices in $G'(\mathcal{P})$. There is no problem at the interfaces between e-edges since we can assume that the z's are to the far right at those interfaces. This is true since the Hamiltonian cycles in $H_{n,k}$ always include

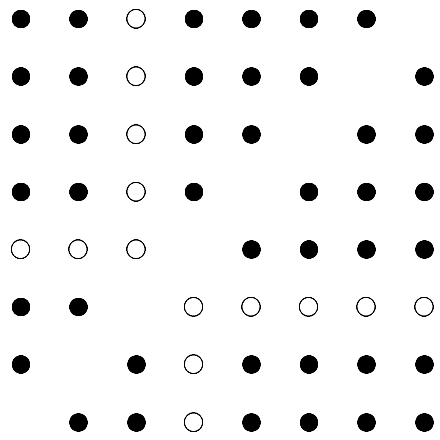

Figure 16: The graph $H_{6,2}$.

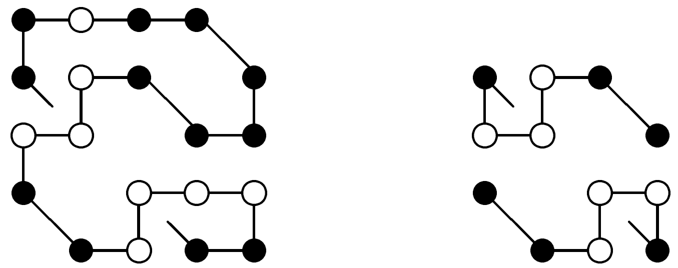

Figure 17: (a) Hamiltonian cycle in $H_{3,1}$, (b) Hamiltonian path in $H_{2,1}$.

the edge e_3. This edge is then deleted to get a Hamiltonian path from Azz to Bzz. $\square$

COROLLARY 1 *The graph $G'(n,m,2)$ has a Hamiltonian path if and only if n and m are both odd (assuming $n, m \geq 1$).*

PROOF: If $n \geq 1$ and $m \geq 1$ are not both odd then the parity difference is greater than one and thus there is no Hamiltonian path. If n and m are both odd then $G'(n,m)$ satisfies the conditions of Lemma 10. $\square$

For the rest of this section we are going to consider what happens when two isolated elements are added to a poset. The resulting poset is always balanced.

DEFINITION 5 A *tree-of-cycles* (TOC) is a graph, defined recursively as follows.

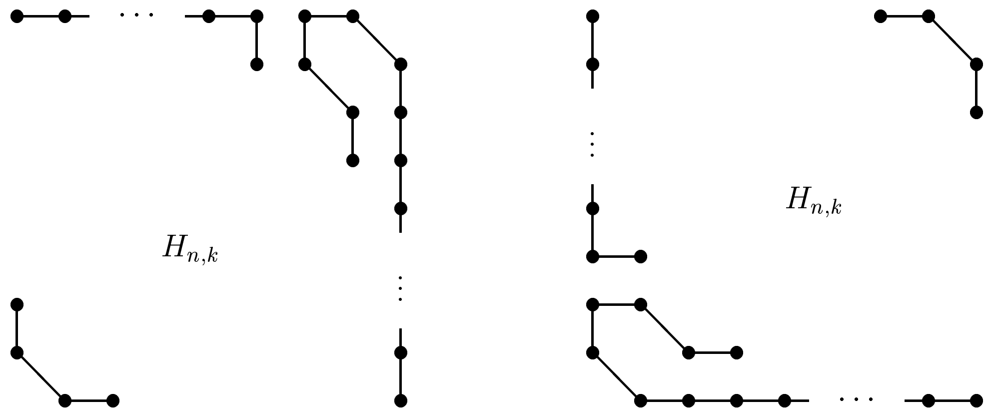

Figure 18: Extending Hamiltonian cycles in $H_{n,k}$ to $H_{n+1,k}$ and $H_{n+1,k+1}$.

1. A single edge is a TOC.
2. An even length cycle is a TOC.
3. If G is a TOC and β is a pendant vertex then the graph G' obtained by adding a new pendant vertex α, adjacent to β, is a TOC.
4. If G is a TOC and β is a pendant vertex then the graph G' obtained by replacing β with an even length cycle is a TOC.
5. If G is a TOC and β is a two-valent vertex on a cycle then the graph G' obtained by adding a new pendant vertex α, adjacent to β, is a TOC.

Another way of defining a TOC is to say that it is a connected graph in which every cycle has even length, every vertex lies on at most one cycle, every vertex on a cycle has degree 2 or 3 and every vertex not on a cycle has degree 1 or 2. We will be interested in TOCs as spanning subgraphs of various adjacent transposition graphs. A spanning TOC is abbreviated STOC.

LEMMA 11 *If $G'(\mathcal{P})$ has a STOC then $G'(\mathcal{P} \cup C_1 \cup C_1)$ has a Hamiltonian cycle.*

PROOF: Denote the C_1 elements by x and y. Each vertex of $G'(\mathcal{P})$ becomes a copy of $G'(1,1,n)$ in $G'(\mathcal{P} \cup C_1 \cup C_1)$, where n is the number of elements of $\mathcal{P}$. If (α, β) is an edge of $G'(\mathcal{P})$ then

$$\begin{aligned} &a = (xy\alpha, xy\beta), \quad b = (\alpha yx, \beta yx), \\ &c = (x\alpha y, x\beta y), \quad d = (y\alpha x, y\beta x) \end{aligned}$$

are all edges in H (as are $(yx\alpha, yx\beta)$, $(\alpha xy, \beta xy)$). In a slight abuse of notation, the vertices of the edges a, b, c, d correspond to the vertices a, b, c, d of $G'(1,1,n)$ as shown in Figure 19. There are two copies of $G'(1,1,n)$; one corresponding to α and the other to β. The proof will be broken into two cases, depending on the parity of n.

Suppose that n is even. Figure 20 shows how to construct a Hamiltonian path between a and b and a Hamiltonian path between c and d. Figure 21 shows how to construct a spanning pair of paths, one path between d and a, the other between b and c (both n odd and n even are shown in the Figure). These paths will be summarized by using the graphs A, B, and C shown in Figure 19. Graphs A and C correspond to the paths of Figure 20 and graph B corresponds to the pair of paths of Figure 21.

Each vertex of $G'(\mathcal{P})$ is replaced by an instance of A,B, or C. If S is the spanning STOC of $G'(\mathcal{P})$ then each pendant vertex of S and each two-valent vertex on a cycle of S are replaced by an A or a C and each two-valent vertex not on a cycle and each degree 3 vertex on a cycle are replaced by a B. We now proceed through the five cases of the recursive definition of a TOC as illustrated in Figure 23. It is helpful to think of the Hamiltonian cycle that we are constructing as being oriented in some direction. This is the reason for the arrows in Figure 23. Consider rule 1. If the edge is (α, β) then the cycle goes from $xy\alpha$ to αyx using a type A path, then the adjacent transposition from α to β is made, then the cycle goes from βyx to $xy\beta$ using a type A path in the opposite direction and finally the adjacent transposition from β to α completes the cycle. The other cases depicted in Figure 23 are interpreted in a similar manner. Note in the case of rules 3, 4, and 5 that the cycle enters β at a and leaves β at b both before and after the TOC is expanded. Each of the replacements depicted in Figure 23 has symmetric versions with the direction of the cycle reversed and/or the roles of A and C reversed.

Now suppose that n is odd. Figure 22 shows how to construct a Hamiltonian path between a and c and a Hamiltonian path between b and d. Figure 21 shows how to construct a spanning pair of paths, one path between d and a, the other between b and c. These paths will be summarized by using the graphs D, E, and B shown in Figure 19. Graphs D and E correspond to the paths of Figure 22 and graph B corresponds to the pair of paths of Figure 21.

Each vertex of $G'(\mathcal{P})$ is replaced by an instance of D, E, or B. If S is the spanning STOC of $G'(\mathcal{P})$ then each pendant vertex of S and each two-valent vertex on a cycle of S are replaced by a D or a E, and each two-valent vertex not on a cycle and each degree 3 vertex on a cycle are replaced by a B. We now proceed through the five cases of the recursive definition of a TOC as illustrated in Figure 24 as was done in the case of n even. $\square$

We now show that $G'(m,n)$ has a STOC. Bitstrings of m 0's and n 1's are used to denote the vertices of $G'(m,n)$. Note that 0^m1^n and 1^n0^m must be pendant vertices in any STOC of $G'(m,n)$ since they are pendant vertices in $G'(m,n)$ itself.

LEMMA 12 *For any $m, n \geq 1$ there is a STOC in $G'(m,n)$. If $m, n \geq 3$ then $G'(m,n)$ has a STOC that contains at least one cycle and pendant vertices* $0^{m-1}1^n0$

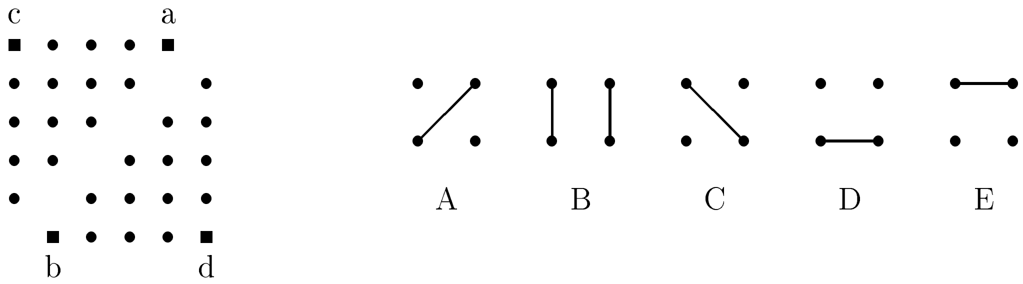

Figure 19: The vertices a, b, c, d and the symbols for the path types A, B, C, D, E.

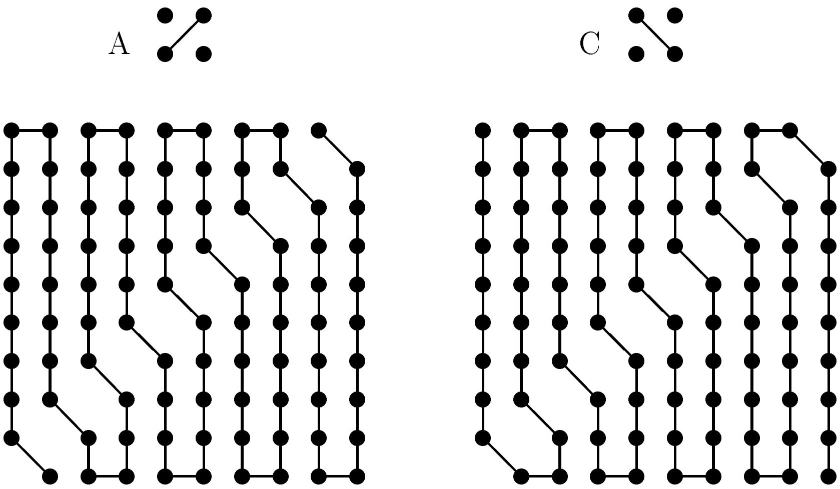

Figure 20: Hamiltonian path from a to b and from c to d (n even).

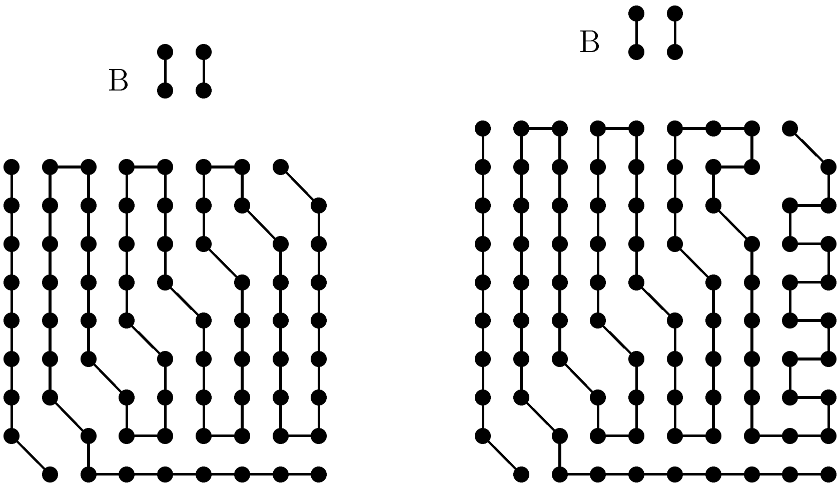

Figure 21: Spanning pair of paths from a to d and from c to b.

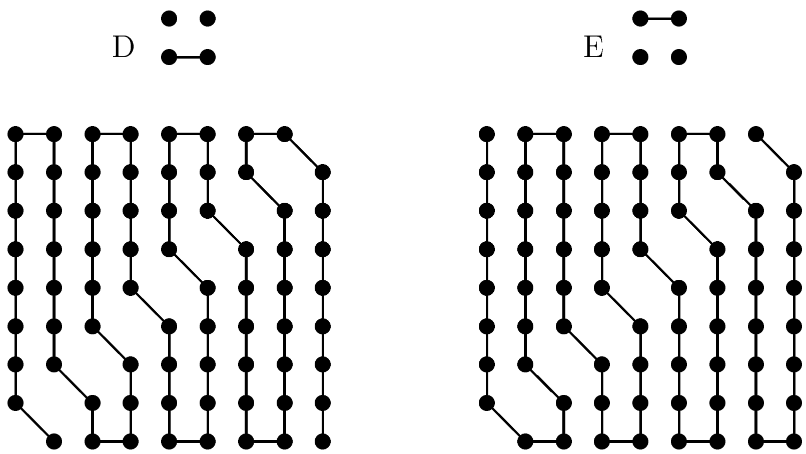

Figure 22: Hamiltonian path from b to d and from a to c (n odd).

and $1^{n-1}0^m1$.

Proof: If $m = 1$ or $n = 1$ then $G'(m, n)$ is a path of length $m + n$. By rules 1 and 3, any path is a STOC. If $m = 2$ and $n \geq 2$ then $G'(m, n)$ is isomorphic to I_n. Figure 25 shows how to construct a STOC in I_n that contains one cycle. Note that the cycle has two-valent vertices 01^n0 and $01^{n-1}01$ (if $n \geq 3$).

If $m = 3$ and $n \geq 3$ then we prove by induction on n that $G'(m, n)$ has a STOC with pendant vertices 001^n0 and $1^{n-1}0001$ and a cycle with two-valent vertex 01^n00. The base case $n = 3$ is shown in Figure 26. If $n > 3$ then we classify the vertices of $G'(3, n)$ according to whether the last bit is a 0 or a 1. That is (in a slight abuse of notation),

$$G'(3, n) \;=\; G'(2, n)0 \cup G'(3, n-1)1$$

Inductively, $G'(3, n-1)1$ has the degree two vertex $a = 01^{n-1}001$ on a cycle in its STOC. By the construction given above for the $m = 2$ case, $G'(2, n)0$ has the degree two vertex $a' = 01^{n-1}010$ on a cycle in its STOC. These two STOCs may be joined by adding the edge $[a, a']$ to get a STOC, S, for $G'(3, n)$. Furthermore, $G'(2, n)0$ has the degree two vertex 01^n00 on a cycle, and thus S does as well.

If $m, n \geq 4$ then we proceed by induction on m and n, using the decomposition of $G'(m, n)$ implicit in the paper of Eades, Hickey and Read [4] (see also Hough and Ruskey [6]) and the $m = 3$ and $n \geq 3$ construction of the previous paragraph as the base case. In [4] the vertex set of $G'(m, n)$ is decomposed into 4 parts, depending on the final two bits of each vertex. That is (again, abusing the notation),

$$G'(m, n) \;=\; G'(m-2, n)00 \cup G'(m-1, n-1)01 \cup G'(m-1, n-1)10 \cup G'(m, n-2)11$$

Inductively, $G'(m-1, n-1)$ has a STOC, S, with a cycle C. Let v and w be two adjacent vertices on C in $G'(m-1, n-1)01$ and let v' and w' be the corresponding vertices in $G'(m-1, n-1)10$. Construct a STOC, call it T, in $G'(m-1, n-1)01 \cup$

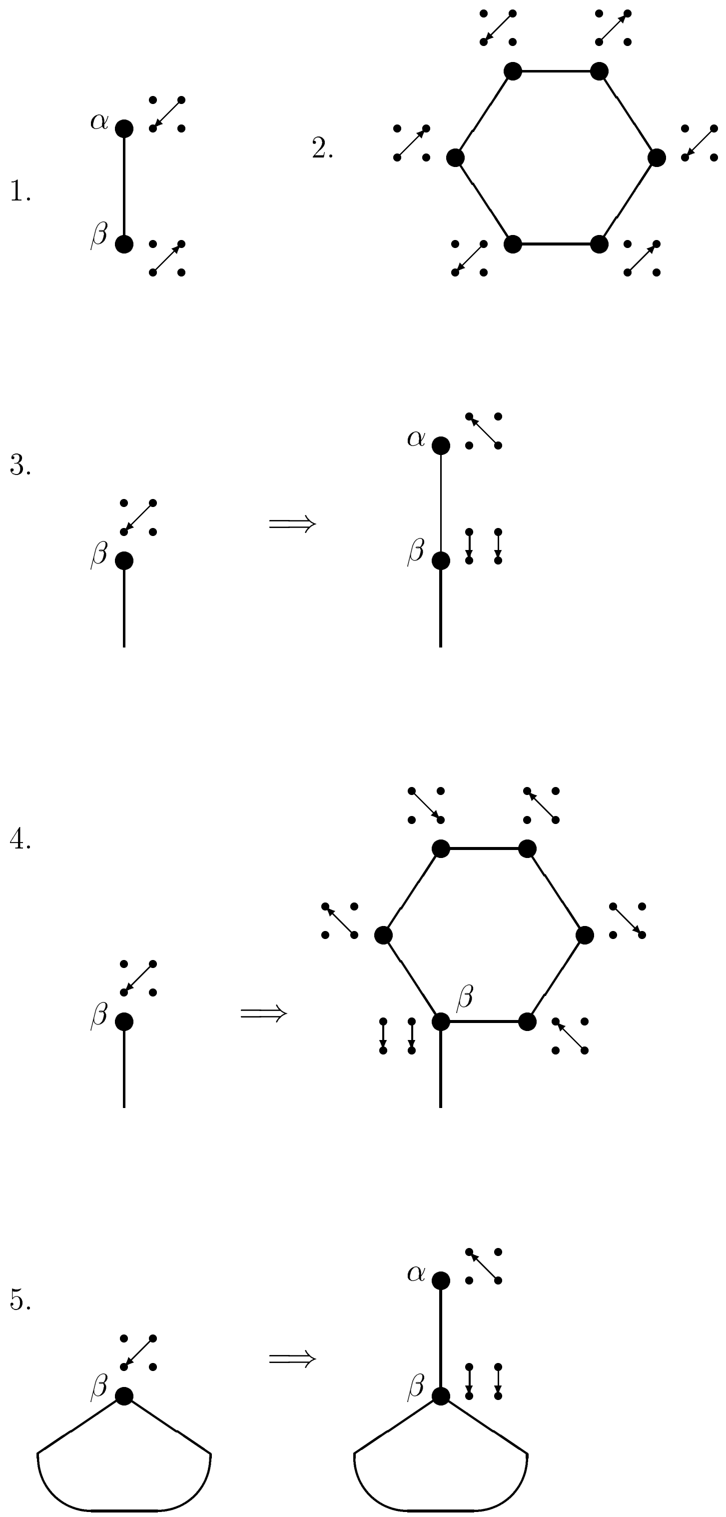

Figure 23: Converting a STOC into a Hamiltonian cycle (n even).

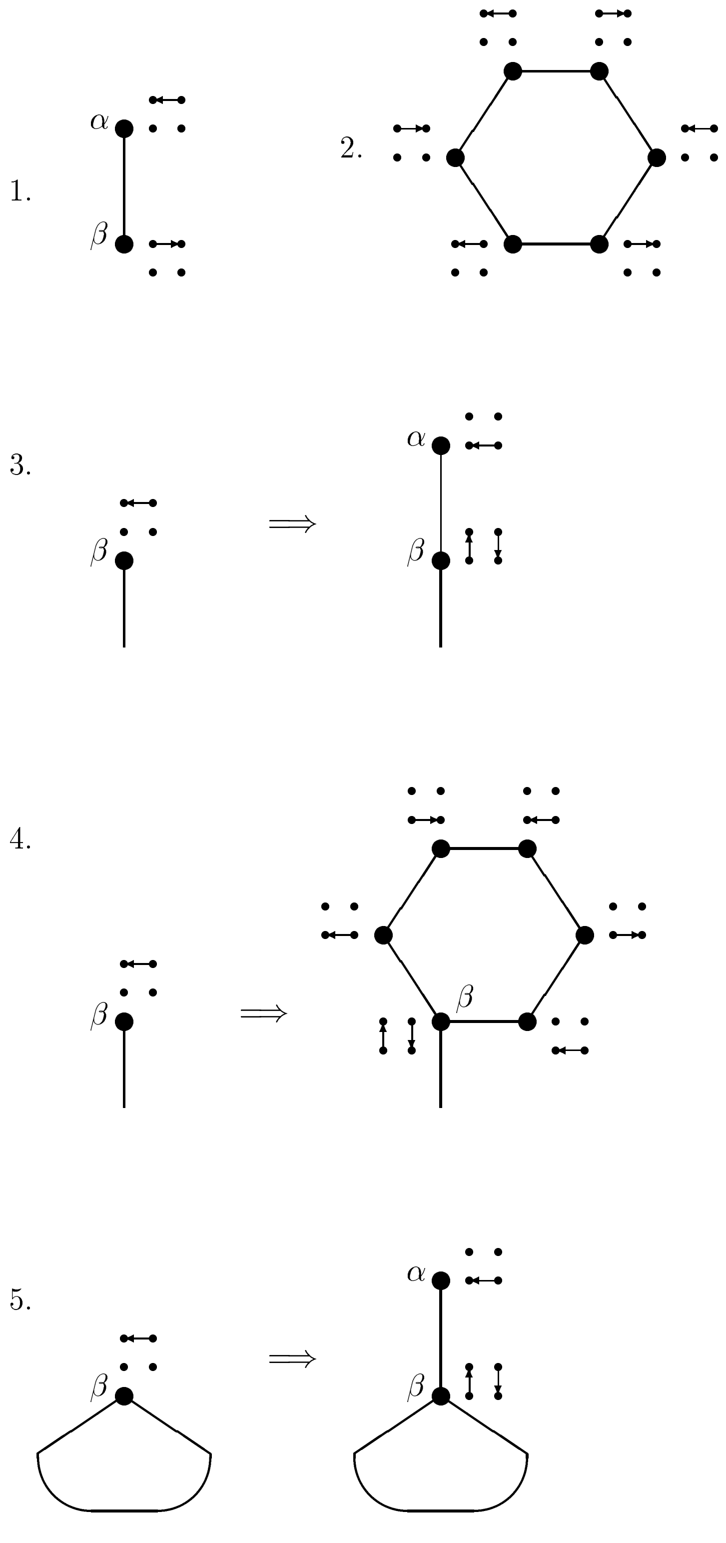

Figure 24: Converting a STOC into a Hamiltonian cycle (n odd).

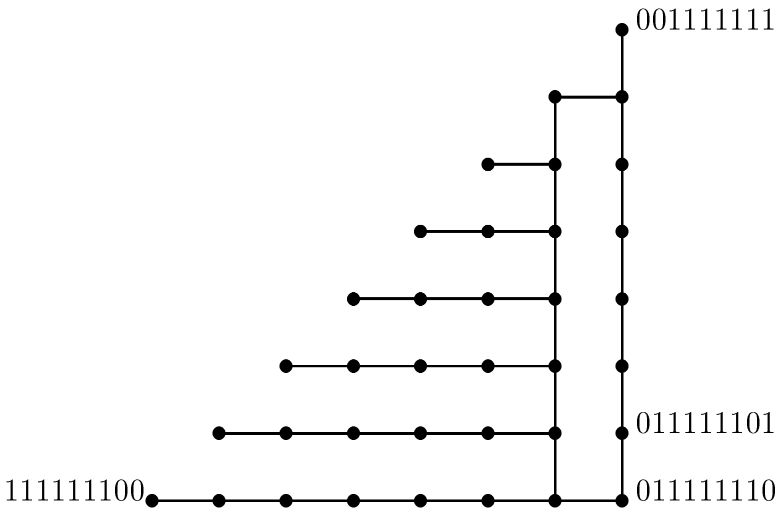

Figure 25: A STOC in I_n (for $n = 7$).

$G'(m-1, n-1)10$ by replacing the edges $[v, w]$ and $[v', w']$ with the edges $[v, v']$ and $[w, w']$. Since S has pendant vertices $0^{m-1}1^{n-1}$ and $1^{n-1}0^{m-1}$, T will have pendant vertices $a = 0^{m-1}1^{n-1}10$ and $a' = 1^{n-1}0^{m-1}01$. Vertices a and a' are the pendant vertices specified in the statement of the lemma.

Inductively, $G'(m-2, n)00$ has a STOC, T_{00}, with pendant vertex $b = 0^{m-2}1^n00$, $G'(m, n-2)11$ has a STOC, T_{11}, with pendant vertex $c = 1^{m-2}0^n11$, and T has a pendant vertices $b' = 0^{m-2}1^{n-1}010$ and $c' = 1^{m-2}0^{n-1}101$. Thus a STOC for all of $G'(m, n)$ may be obtained by adding the edge $[b, b']$ to join T and T_{00} and adding the edge $[c, c']$ to join T and T_{11}. □

THEOREM 4 *The graph $G'(1, 1, m, n)$ has a Hamiltonian cycle for any $m, n \geq 1$.*

PROOF: This follow directly from Lemmas 11 and 12. □

4.3 Counting Edges

Although not of use in the discussion above, it is of some interest to know the number of edges in the adjacent transposition graphs. Defining $f(\mathcal{P})$ to be the number of edges in $G'(\mathcal{P})$, the following recurrence relation holds.

$$f(\mathcal{P}) = \sum_{x \in Min(\mathcal{P})} f(\mathcal{P} - x) + \sum_{x,y \in Min(\mathcal{P})} e(\mathcal{P} - x - y) \qquad (6)$$

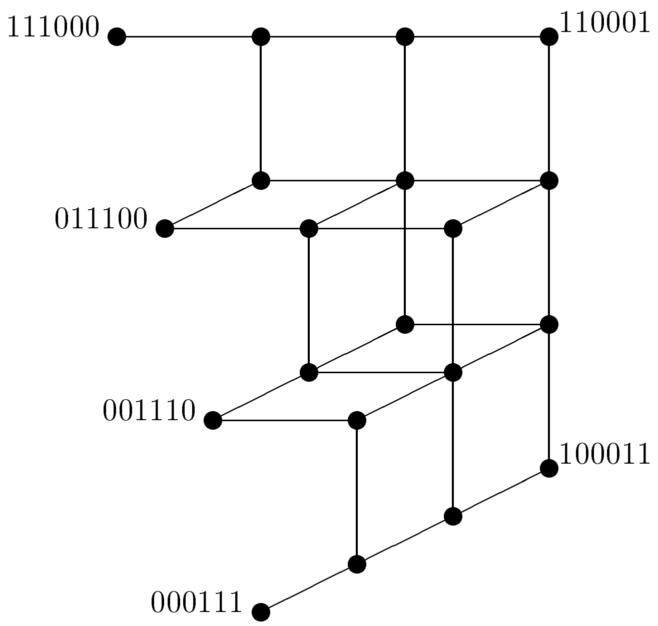

Figure 26: A STOC in $G'(3,3)$.

This recurrence relation can be solved in the case of multisets. The number of edges in $G'(n_0, n_1, \ldots, n_t)$ is

$$\begin{aligned} f(\mathbf{n}) &= (n-1) \sum_{0 \le i < j \le t} \binom{n-2}{n_0, \ldots, n_i - 1, \ldots, n_j - 1, \ldots, n_t} \\ &= \frac{1}{2}(n-1) \left[\binom{n}{n_0, \ldots, n_t} - \sum_{0 \le i \le t} \binom{n-2}{n_0, \ldots, n_i - 2, \ldots, n_t} \right] \end{aligned}$$

The first expression is obtained by counting the number of edges induced by one of the $n-1$ adjacent pairs of *positions* within a permutation. The last expression simplifies for unrestricted permutations (an antichain) and for combinations (two disjoint chains).

$$\text{Permutations: } f(1, 1, \ldots, 1) = \frac{1}{2}(n-1)n!$$

$$\text{Combinations: } f(m, n) = \frac{(m+n-1)!}{(m-1)!(n-1)!}$$

It would be nice to have something analogous to (5) for counting the number of edges in the case where the poset is a tree.

5 Final Remarks

We have raised more questions than have been answered. The main problem that we would like to see resolved is the following conjecture:

Conjecture 1 *If $\mathcal{P}$ is a balanced poset then $G(\mathcal{P})$ has a Hamiltonian path.*

This conjecture is known to be true for posets of disjoint chains (this paper), for the binary tree posets (e.g., Figure 2, see [20]), for fences ([19]), for crowns, subset lattices, partition lattices, etc. ([14]), and for many other small posets. In fact, evidence indicates that the conjecture may even be true with G replaced by G'. The following questions will be the subject of future investigations.

1. Does a Hamiltonian path exist in $G'(\mathbf{n})$ whenever $d(\mathbf{n})$ is zero?

2. Does a Hamiltonian *cycle* exist in $G(\mathbf{n})$ whenever $d(\mathbf{n})$ is zero and $t > 1$?

3. Does a Hamiltonian cycle exist in $G'(\mathbf{n})$ whenever $d(\mathbf{n})$ is zero and $\mathbf{n}$ is neither of n, m with n and m both odd or $1, 1, m$ with m odd?

4. If n and m are odd, then does $G(n, m) - \{0^n 1^m, 1^m 0^n\}$ have a Hamiltonian cycle? In other words, are the two pendant vertices the "only" reason that $G(n, m)$ is not Hamiltonian?

5. What is $d(\mathcal{P})$ if $\mathcal{P}$ is an order ideal $\mathcal{I}$ of $C_m \times C_n$? The linear extensions of $\mathcal{I}$ correspond to Young Tableau of shape specified by $\mathcal{I}$.

6. What is $d(C_n \times C_m)$ and does $G(C_m \times C_n)$ have a Hamiltonian path whenever its parity difference is zero? The Hasse diagram of $C_m \times C_n$ is an m by n grid rotated by 45°. We can show that $d(C_n \times C_m) = 0$ whenever m and n are both even by using a parity reversing involution ϕ on the set of linear extensions. Consider a linear extension

$$\pi = a_1 a_2 \cdots a_{mn}$$

Let a_{2i}, a_{2i+1} be the leftmost pair of incomparable elements (a_1 ignored since it is the minimum). If there is no such pair then the involution returns π (i.e., π is a fixed-point), and if there is such a pair then

$$\phi(\pi) = a_1 a_2 \cdots a_{2i-1} a_{2i+1} a_{2i} a_{2i+2} \cdots a_{mn}$$

If n and m are both even, then any fixed point induces a perfect matching in the Hasse diagram of $C_m \times C_n$ with the maximum and minimum elements removed. However, there is no perfect matching in an m by n grid with two extreme vertices removed. The reasoning is the same as that used to solve the well-known problem of covering a chess board with two opposite corners removed with dominoes.

If m and n are both odd, then ϕ does have fixed points, but the parity difference seems to be zero (see table below).

If $m = 2$ then the linear extensions correspond to binary trees, and it is shown by Ruskey and Proskurowski [20] and Ruskey [16] that the parity difference is

$$d(C_2 \times C_n) = \begin{cases} e(C_2 \times C_{(n-1)/2}) & \text{if } n \text{ odd} \\ 0 & \text{if } n \text{ is even} \end{cases} \tag{7}$$

The number $e(C_2 \times C_n)$ is the nth Catalan number $(2n)!/((n+1)!n!)$. Equation (7) can be derived using the involution ϕ given above. If n is odd then every fixed point induces the same perfect matching in the grid with two extreme vertices removed and thus the fixed points all have the same parity. Collapsing the poset $C_2 \times C_n$, as was done in Lemma 3, yields the poset $C_2 \times C_{(n-1)/2}$.

The following table was generated by computer.

	1	2	3	4	5	6
1	1					
2	1	0				
3	1	1	0			
4	1	0	2	0		
5	1	2	0	12	0	
6	1	0	12	0	286	0
7	1	5	0	286	0	33592
8	1	0	110	0	33592	0
9	1	14	0	12376	?	?
10	1	0	1274	0	?	0
11	1	42	0	?	?	?
12	1	0	17136	0	?	0
13	1	132	0	?	?	?
14	1	0	255816	0	?	0

7. An easier question than The Conjecture is whether $G(\mathcal{P})$ has a perfect matching for every balanced poset $\mathcal{P}$.

8. If the parity difference is not zero, then a Hamiltonian cycle does not exist in the transposition graph. To pose an interesting question in the case of $d(\mathcal{P}) > 0$, we may ask for the shortest spanning closed walk. The length (number of edges) in a spanning closed walk is at least $e(\mathcal{P}) + d(\mathcal{P})$; a spanning closed walk with this length is said to be *minimal*. Lehmer [13] considered a special type of spanning closed walk where any vertex occurs at most twice and those that occur twice

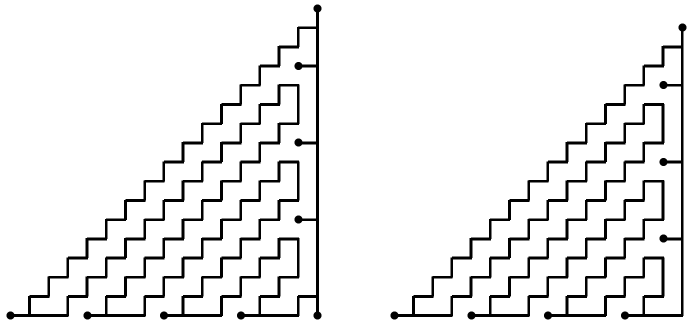

Figure 27: (a) $n \equiv 0 \bmod 4$ (b) $n \equiv 3 \bmod 4$

are separated by only one vertex. Such a walk can be modeled by what we call a *cyclic comb*. This is a cycle with some pendant edges. More precisely, a cyclic comb is a connected unicyclic graph in which every vertex not on the cycle has degree one and every vertex on the cycle has degree two or three. There is a minimal cyclic comb in $G'(2, n)$ for any $n \geq 2$. The exact comb depends on $n \bmod 4$. See Figure 27. The strongest statement to which we do not have a counterexample is: $G'(\mathcal{P})$ has a cyclic comb with $d(\mathcal{P})$ pendant vertices.

6 Acknowledgement.

The presentation of this paper has been greatly improved as a result of numerous detailed comments and suggestions by one of the referees. The author would also like to thank Michelle Wachs, Donald J. Miller and Dominique Roelants van Baronaigien for helpful discussions and Ken Wong for writing some programs.

References

[1] J.A. Bondy and U.S.R. Murty. *Graph Theory with Applications*. North-Holland, 1976.

[2] M. Buck and D. Wiedemann. Gray codes with restricted density. *Discrete Math.*, 48:163–171, 1984.

[3] M. Chein and M. Habib. The jump number of dags and posets: an introduction. *Discrete Applied Math.*, 9:189–194, 1980.

[4] P. Eades, M. Hickey, and R.C. Read. Some Hamilton paths and a minimal change algorithm. *JACM*, 31:19–29, 1984.

[5] P.C. Fishburn and W.V. Gehrlein. Minimizing bumps in linear extensions of ordered sets. *Order*, 1:3–14, 1986.

[6] T. Hough and F. Ruskey. An efficient implementation of the Eades, Hickey, Read adjacent interchange combination generation algorithm. *J. Combinat. Math. and Combinat. Computing*, 4:79–86, 1988.

[7] S.M. Johnson. Generation of permutations by adjacent transpositions. *Math. Computation*, 17:282–285, 1963.

[8] A.D. Kalvin and Y.L. Varol. On the generation of all topological sortings. *J. Algorithms*, 4:150–162, 1983.

[9] D.E. Knuth. *Sorting and Searching.* Addison-Wesley, 1973.

[10] D.E. Knuth. Lexicographic permutations with restrictions. *Discrete Applied Math.*, 1:117–125, 1979.

[11] D.E. Knuth and J. Szwarcfiter. A structured program to generate all topological sorting arrangements. *Information Processing Letters*, 2:153–157, 1974.

[12] C.W. Ko and F. Ruskey. Solution of some multi-dimensional lattice path parity difference recurrence relations. *Discrete Math.*, 71:47–56, 1988.

[13] D.H. Lehmer. Permutation by adjacent interchanges. *American Math. Monthly*, 72:36–46, 1965.

[14] G. Pruesse and F. Ruskey. Generating the linear extensions of certain posets by transpositions. *SIAM J. Discrete Mathematics*, to appear.

[15] I. Rival. Optimal linear extensions by interchanging chains. *Proc. American Math. Society*, 89:387–394, 1983.

[16] F. Ruskey. Solution of some lattice path parity difference recurrence relations using involutions. *Congressus Numerantium*, 59:257–266, 1987.

[17] F. Ruskey. Adjacent interchange generation of combinations. *J. Algorithms*, 9:162–180, 1988.

[18] F. Ruskey. Research problem 90. *Discrete Math.*, 70:111–112, 1988.

[19] F. Ruskey. Transposition generation of alternating permutations. *Order*, 6:227–233, 1989.

[20] F. Ruskey and A. Proskurowski. Generating binary trees by transpositions. *J. Algorithms*, 11:68–84, 1990.

[21] H. Steinhaus. *One Hundred Problems in Elementary Mathematics.* Basic, 1964.

[22] M.M. Sysło. A graph-theoretic approach to the jump number problem. In I. Rival, editor, *Graphs and Order*, pages 185–215. Reidel, Dordrecht, 1985.

[23] H.F. Trotter. Algorithm 115: Perm. *Comm. ACM*, 5:434–435, 1962.

[24] Y.L. Varol and D. Rotem. An algorithm to generate all topological sorting arrangements. *Computer J.*, 24:83–84, 1981.

[25] N. Zaguia. Minimizing bumps for posets of width two. *Discrete Applied Math.*, 21:157–162, 1988.