

FRIENDS AND STRANGERS WALKING ON GRAPHS

Colin Defant*¹ and Noah Kravitz²

¹*Department of Mathematics, Princeton University, Princeton, NJ 08540, U.S.A.
cdefant@princeton.edu*

²*Department of Mathematics, Princeton University, Princeton, NJ 08540, U.S.A.
nkravitz@princeton.edu*

Submitted: Sep 17, 2020; Accepted: Jun 14, 2021; Published: Dec 15, 2021

© The authors. Released under the CC BY license (International 4.0).

Abstract. Given graphs X and Y with vertex sets $V(X)$ and $V(Y)$ of the same cardinality, we define a graph $\text{FS}(X, Y)$ whose vertex set consists of all bijections $\sigma: V(X) \rightarrow V(Y)$, where two bijections σ and σ' are adjacent if they agree everywhere except for two adjacent vertices $a, b \in V(X)$ such that $\sigma(a)$ and $\sigma(b)$ are adjacent in Y . This setup, which has a natural interpretation in terms of friends and strangers walking on graphs, provides a common generalization of Cayley graphs of symmetric groups generated by transpositions, the famous 15-puzzle, generalizations of the 15-puzzle as studied by Wilson, and work of Stanley related to flag h -vectors. We derive several general results about the graphs $\text{FS}(X, Y)$ before focusing our attention on some specific choices of X . When X is a path graph, we show that the connected components of $\text{FS}(X, Y)$ correspond to the acyclic orientations of the complement of Y . When X is a cycle, we obtain a full description of the connected components of $\text{FS}(X, Y)$ in terms of toric acyclic orientations of the complement of Y . We then derive various necessary and/or sufficient conditions on the graphs X and Y that guarantee the connectedness of $\text{FS}(X, Y)$. Finally, we raise several promising further questions.

Mathematics Subject Classifications. 05C40, 05C38, 05A05

1. Introduction

Let X be a simple graph with n vertices. Imagine that n different people, any two of whom are either friends or strangers, are standing so that one person is at each vertex of X . At each point in time, two friends standing at adjacent vertices of X may swap places by simultaneously crossing the edge that connects their respective vertices; two strangers may not swap places in this way. Our goal is to understand which configurations of people can be reached from other configurations when we allow the people to swap places repeatedly in this manner. The

*Supported by a Fannie and John Hertz Foundation Fellowship and an NSF Graduate Research Fellowship (grant number DGE-1656466).

resolution of this problem certainly depends on the graph X , as well as on the information about which people are friends with each other; this further information can be encoded in a graph Y whose vertex set is the set of people and where edges indicate friendship.

To formalize and refine this problem, we define the *friends-and-strangers graph* $\text{FS}(X, Y)$ whose vertex set is the set of bijections $\sigma: V(X) \rightarrow V(Y)$. Two bijections $\sigma, \sigma': V(X) \rightarrow V(Y)$ are adjacent in $\text{FS}(X, Y)$ if and only if we can find an edge $\{a, b\}$ in X such that:

- $\{\sigma(a), \sigma(b)\}$ is an edge in Y ;
- $\sigma(a) = \sigma'(b)$ and $\sigma(b) = \sigma'(a)$;
- $\sigma(c) = \sigma'(c)$ for all $c \in V(X) \setminus \{a, b\}$.

When this is the case, we refer to the operation that transforms σ into σ' as an *(X, Y) -friendly swap across $\{a, b\}$* . Performing an (X, Y) -friendly swap corresponds to allowing two friends to swap places in the graph X . Notice that the isomorphism type of $\text{FS}(X, Y)$ depends only on the isomorphism types of X and Y . Since we will usually be concerned only with the graph-theoretic structure of $\text{FS}(X, Y)$ (such as the number and sizes of connected components), we will often specify the graphs X and Y only up to isomorphism.

It is sometimes convenient to assume that $V(X)$ and $V(Y)$ are both the set $[n] := \{1, \dots, n\}$. In this case, the vertices of $\text{FS}(X, Y)$ are the elements of the symmetric group \mathfrak{S}_n , which consists of all permutations of the numbers $1, \dots, n$. For $i, j \in [n]$, let (ij) be the transposition in \mathfrak{S}_n that swaps the numbers i and j . If $\sigma \in \mathfrak{S}_n$ is such that $\{i, j\}$ is an edge in X and $\{\sigma(i), \sigma(j)\}$ is an edge in Y , then we can perform an (X, Y) -friendly swap across $\{i, j\}$ to change σ into the permutation $\sigma \circ (ij)$. If we write the permutation σ in one-line notation as $\sigma = \sigma(1) \dots \sigma(n)$, then an (X, Y) -friendly swap transposes two entries of the permutation such that the positions of the entries are adjacent in X and the entries themselves are adjacent in Y .

Example 1.1. If $X = Y = \begin{array}{c} \bullet \\ \text{---} \\ \bullet \text{---} \bullet \\ \text{---} \\ \bullet \end{array}$, then $\text{FS}(X, Y) = \begin{array}{c} \bullet \text{---} \bullet \text{---} \bullet \\ \text{---} \\ \bullet \text{---} \bullet \text{---} \bullet \\ \text{---} \\ \bullet \end{array}$.

Example 1.2. If $X = \begin{array}{c} \bullet \\ \diagdown \quad \diagup \\ \bullet \text{---} \bullet \\ \diagup \quad \diagdown \\ \bullet \end{array}$ and $Y = \begin{array}{c} \bullet \\ \diagdown \quad \diagup \\ \bullet \text{---} \bullet \\ \diagup \quad \diagdown \\ \bullet \end{array} \text{---} \bullet$, then $\text{FS}(X, Y)$ is the graph

shown (without vertex labels) in Figure 1.1. Notice that $\text{FS}(X, Y)$ has 5 pairwise isomorphic connected components, each of which exhibits a highly symmetric structure; our results in Sections 2 and 4 will explain both of these properties.

In this paper, we concern ourselves mainly with the problem of understanding the structure of $\text{FS}(X, Y)$ for various graphs X and Y . This is a very general setup, of which some previously-considered problems are special cases. One motivation for phrasing the problem in terms of permutations is to encompass previous work on multiplication by restricted transpositions. To give some examples, we let K_n denote the complete graph on the vertex set $[n]$ and let Path_n be the path graph on the vertex set $[n]$ with edges $\{i, i+1\}$ for all $i \in [n-1]$. The graph $\text{FS}(K_n, K_n)$




Figure 1.1: The graph $FS(X, Y)$ from Example 1.2.

is the Cayley graph of \mathfrak{S}_n generated by all transpositions. The graph $FS(K_n, Path_n)$ is the weak Bruhat graph of \mathfrak{S}_n , which is also the 1-skeleton of the standard $(n - 1)$ -dimensional permutohedron. More generally, if Y is connected, then $FS(K_n, Y)$ is the Cayley graph of \mathfrak{S}_n generated by the transpositions corresponding to the edges of Y . Letting \bar{Y} denote the complement of Y , we find that $FS(Path_n, \bar{Y})$ is the same as the update graph introduced by Barrett and Reidys in their work [2] on sequential dynamical systems. Stanley’s article [22] focuses on the connected components of $FS(Path_n, Path_n)$.

Another special case of our setup has previously been phrased in terms of sliding tiles on graphs. Analyzing the famous 15-puzzle is equivalent to analyzing the connected components of $FS(Grid_{4 \times 4}, Star_{16})$, where $Grid_{4 \times 4}$ is a 4×4 grid graph and $Star_n$ denotes the star graph with vertex set $[n]$ and edges $\{i, n\}$ for all $i \in [n - 1]$. The purpose of Wilson’s article [23] is to generalize the 15-puzzle by computing the number of connected components of $FS(X, Star_n)$ for arbitrary graphs X . Conway, Elkies, and Martin [9] investigated a “projective” variant of the 15-puzzle problem and its connections to the Mathieu groups. More recently, Yang [24] considered a similar problem in this spirit. See also [4, 13].

A simple but useful observation is that for any n -vertex simple graphs X and Y , the graph $FS(X, Y)$ is isomorphic to $FS(Y, X)$. Indeed, the map $\sigma \mapsto \sigma^{-1}$ provides the necessary isomorphism. For example, Section 3 concerns the graphs of the form $FS(Path_n, Y)$; by taking inverses of permutations, one could easily rewrite the results in that section in terms of graphs of the form $FS(X, Path_n)$.

In Section 2, we prove several general properties of the graphs $FS(X, Y)$. For example, we show how the decomposition of X into connected components translates into a decomposition of $FS(X, Y)$ into a disjoint union of Cartesian products. We prove that $FS(X, Y)$ is disconnected whenever X and Y both have cut vertices, and we provide a lower bound for the number of connected components. We also show that $FS(X, Y)$ is disconnected whenever X and Y are both bipartite.

In Section 3, we investigate $FS(Path_n, Y)$ for arbitrary graphs Y . This generalizes Stanley’s work in [22]. We will see that the connected components correspond bijectively to the acyclic orientations of the complement of Y and that the vertices in each connected component are the

linear extensions of a poset naturally associated to the corresponding acyclic orientation. This result is closely related to a known result concerning Coxeter elements in Coxeter systems. We will also see that each connected component of $\text{FS}(\text{Path}_n, Y)$ is isomorphic to the Hasse diagram of a convex subset of the weak order on \mathfrak{S}_n .

The characterization of the connected components of $\text{FS}(\text{Path}_n, Y)$ is also roughly equivalent to a simple yet useful fact about linear extensions of posets, which we state below as Proposition 1.3. Let $P = ([n], \leq_P)$ be a poset with underlying set $[n]$. A *linear extension* of P is a permutation $\sigma \in \mathfrak{S}_n$ such that $\sigma^{-1}(a) \leq \sigma^{-1}(b)$ whenever $a \leq_P b$. Let $\mathcal{L}(P)$ denote the set of linear extensions of P . For each $i \in [n-1]$, define the *toggle operator* $t_i: \mathcal{L}(P) \rightarrow \mathcal{L}(P)$ by

$$t_i(\sigma) = \begin{cases} \sigma \circ (i \ i+1), & \text{if } \sigma(i) \not\leq_P \sigma(i+1); \\ \sigma, & \text{if } \sigma(i) \leq_P \sigma(i+1). \end{cases}$$

It is important to note that for each $i \in [n]$ and $\sigma \in \mathcal{L}(P)$, the permutation $t_i(\sigma)$ is in $\mathcal{L}(P)$. Furthermore, the map $t_i: \mathcal{L}(P) \rightarrow \mathcal{L}(P)$ is an involution.

Proposition 1.3. *Let $P = ([n], \leq_P)$ be a poset. If $\sigma, \sigma' \in \mathcal{L}(P)$, then there exists a sequence i_1, \dots, i_ℓ of elements of $[n-1]$ such that $\sigma' = (t_{i_\ell} \circ \dots \circ t_{i_1})(\sigma)$.*

Proposition 1.3 is not too difficult to prove by induction on n , but it is of fundamental importance in many different places in combinatorics. It is usually used to show that a definition involving a choice of a linear extension is well-defined in the sense that it does not depend on the choice of the linear extension. For example, it has been employed in order to understand Coxeter elements in Coxeter systems [10, Theorem 1.15], Jucys–Murphy elements and Gelfand–Tsetlin bases [17, Proposition 8.18], \oplus -diagrams and \lrcorner -diagrams [15, Proposition 2.6], and the hook-length formula [19, Lemma 3].

In Section 4, we completely describe the connected components of $\text{FS}(\text{Cycle}_n, Y)$, where Cycle_n is the cycle graph with vertex set $[n]$ and edge set $\{\{i, i+1\} : 1 \leq i \leq n-1\} \cup \{\{n, 1\}\}$. This description is much more involved than (yet very similar in flavor to) the description of the connected components of $\text{FS}(\text{Path}_n, Y)$; it makes use of toric acyclic orientations, which have appeared in many contexts and were formalized in [10]. The connected components of $\text{FS}(\text{Cycle}_n, Y)$ can be understood via a new equivalence relation on acyclic orientations of the complement of Y , which we call *double-flip equivalence*. This new notion could be of independent interest; it turns out that our analysis of the graphs $\text{FS}(\text{Cycle}_n, Y)$ not only requires an understanding of double-flip equivalence classes but also reciprocally yields interesting structural information about the double-flip equivalence classes. We will see that each toric acyclic orientation of the complement of Y corresponds to ν isomorphic connected components of $\text{FS}(\text{Cycle}_n, Y)$, where ν is the greatest common divisor of the sizes of the connected components of the complement of Y . One corollary is that $\text{FS}(\text{Cycle}_n, Y)$ is connected if and only if the complement of Y is a forest whose trees have (not necessarily pairwise) coprime sizes. Since Proposition 1.3 is equivalent to our characterization of the connected components of $\text{FS}(\text{Path}_n, Y)$, one can view our characterization of the connected components of $\text{FS}(\text{Cycle}_n, Y)$ as providing a toric analogue of Proposition 1.3.

Section 5 concerns a sufficient condition for $\text{FS}(X, Y)$ to be connected. We phrase this result in terms of hereditary classes and Hamiltonian paths. More precisely, we define a *prolongation*

of a graph X to be a graph \tilde{X} with a Hamiltonian path that itself contains a Hamiltonian path of a subgraph $X^\#$ of \tilde{X} such that $X^\#$ is isomorphic to X . We will prove that if X has a Hamiltonian path and $\text{FS}(X, Y)$ is connected whenever Y belongs to a hereditary class \mathcal{H} , then $\text{FS}(\tilde{X}, \tilde{Y})$ is connected whenever \tilde{X} is a prolongation of X and \tilde{Y} is in \mathcal{H} . As a corollary, we produce several infinite classes of pairs (X, Y) such that $\text{FS}(X, Y)$ is connected.

In Section 6, we turn our attention to necessary conditions for $\text{FS}(X, Y)$ to be connected. We will prove that if X has an induced subgraph isomorphic to Path_d such that the internal vertices of the path all have degree 2 in X and $\text{FS}(X, Y)$ is connected, then the minimum degree of Y is at least $d + 1$. By combining this with the main result from Section 5, we will see that $\text{FS}(\text{Lollipop}_{n-3,3}, Y)$ is connected if and only if Y has minimum degree at least $n - 2$, where $\text{Lollipop}_{n-3,3}$ is a lollipop graph (which is obtained by identifying an endpoint of a path on $n - 2$ vertices with a vertex in the triangle K_3). We will also see that $\text{FS}(\text{Lollipop}_{n-3,3}, Y)$ is connected if and only if $\text{FS}(D_n, Y)$ is connected, where D_n is the Dynkin diagram of type D_n (which is obtained from Path_{n-1} by adding the vertex n and the edge $\{n - 2, n\}$).

We end the paper with numerous suggestions for future work in Section 7. In particular, we define a generalization of friends-and-strangers graphs in which symmetric groups are replaced by arbitrary groups; this construction is given by the intersection of right and left Cayley graphs of the group. Let us also remark that it is very natural to ask probabilistic and extremal questions concerning friends-and-strangers graphs; along with Noga Alon, we address such questions in the article [1].

1.1. Notation and terminology

Throughout this article, we assume all graphs are simple. Let $V(G)$ and $E(G)$ denote the vertex set and edge set of a graph G . Some specific graphs with vertex set $[n]$ that will play a large role for us are:

- the *complete graph* K_n , which has edge set $E(K_n) = \{\{i, j\} : 1 \leq i < j \leq n\}$;
- the *star graph* Star_n , which has edge set $E(\text{Star}_n) = \{\{i, n\} : i \in [n - 1]\}$;
- the *path graph* Path_n , which has edge set $E(\text{Path}_n) = \{\{i, i + 1\} : i \in [n - 1]\}$;
- the *cycle graph* Cycle_n , which has edge set $E(\text{Cycle}_n) = \{\{i, i + 1\} : i \in [n - 1]\} \cup \{\{n, 1\}\}$.

The *disjoint union* of two graphs G_1, G_2 , denoted $G_1 \oplus G_2$, is the graph whose vertex set is the disjoint union $V(G_1) \sqcup V(G_2)$ and whose edge set is the disjoint union $E(G_1) \sqcup E(G_2)$. This definition readily extends to the disjoint union of a family of graphs G_i for i in an index set I ; we denote the resulting disjoint union by $\bigoplus_{i \in I} G_i$. The *Cartesian product* of the graphs G_1, \dots, G_r , written $G_1 \square \dots \square G_r$, has vertex set given by the set-theoretic Cartesian product $V(G_1) \times \dots \times V(G_r)$, where the vertices (a_1, \dots, a_r) and (b_1, \dots, b_r) are adjacent if and only if there is an index $i \in [r]$ such that $\{a_i, b_i\} \in E(G_i)$ and $a_j = b_j$ for all $j \in [r] \setminus \{i\}$.

Some additional terminology concerning a graph G is as follows:

- The *complement* of G , denoted \overline{G} , is the graph with vertex set $V(\overline{G}) = V(G)$ such that for all $a, b \in V(G)$ with $a \neq b$, we have $\{a, b\} \in E(\overline{G})$ if and only if $\{a, b\} \notin E(G)$.
- An *isomorphism* from G to a graph G' is a bijection $\varphi: V(G) \rightarrow V(G')$ such that $\{a, b\} \in E(G)$ if and only if $\{\varphi(a), \varphi(b)\} \in E(G')$. If such an isomorphism exists, we say G and G' are *isomorphic*, denoted $G \cong G'$. An *automorphism* of G is an isomorphism from G to itself.
- We say a graph H is a *subgraph* of G if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$; we say the subgraph H is *induced* if $E(H) = \{\{a, b\} \in E(G) : a, b \in V(H)\}$. Given a subset V_0 of $V(G)$, let $G|_{V_0}$ denote the induced subgraph of G with vertex set V_0 .
- We say G is *bipartite* if there exists a partition of $V(G)$ into two nonempty sets A and B such that every edge in $E(G)$ has one endpoint in A and one endpoint in B ; in this case, the pair $\{A, B\}$ is called a *bipartition* of G .
- We say G is *connected* if for all $a, b \in V(G)$, there is a path in G connecting a to b ; a *connected component* of G is a maximal connected subgraph of G . The *size* of a connected component H of G is $|V(H)|$. Notice that if H_1, \dots, H_r are the connected components of G , then $G = \bigoplus_{i=1}^r H_i$.
- A *cut vertex* of G is a vertex $v \in V(G)$ such that the graph $G|_{V(G) \setminus \{v\}}$ obtained by deleting v has more connected components than G . We say G is *separable* if it is disconnected or has a cut vertex. We say G is *biconnected* if it has at least 2 vertices and is not separable.
- If G is a graph on n vertices, then a *Hamiltonian path* in G is a subgraph of G isomorphic to Path_n .

2. General Properties of the Graphs $\text{FS}(X, Y)$

To initiate the investigation of the graphs $\text{FS}(X, Y)$, we list some general properties that hold when X and Y are taken from broad classes of graphs. Some of these properties are nicely illustrated by Wilson's results for the graphs of the form $\text{FS}(\text{Star}_n, Y)$, so we will recall those results in this section as well. Several of the general results we prove here will be useful when we consider more specific families of graphs in subsequent sections.

Proposition 2.1. *Let $X, \tilde{X}, Y, \tilde{Y}$ be graphs on n vertices. If X is isomorphic to a subgraph of \tilde{X} and Y is isomorphic to a subgraph of \tilde{Y} , then $\text{FS}(X, Y)$ is isomorphic to a subgraph of $\text{FS}(\tilde{X}, \tilde{Y})$.*

Taking $\tilde{X} = \tilde{Y} = K_n$ in this proposition shows that $\text{FS}(X, Y)$ is isomorphic to a subgraph of $\text{FS}(K_n, K_n)$, the Cayley graph of \mathfrak{S}_n generated by the set of all transpositions in \mathfrak{S}_n .

Proof. We may assume that $V(X) = V(\tilde{X}) = V(Y) = V(\tilde{Y}) = [n]$ and that X and Y are subgraphs of \tilde{X} and \tilde{Y} , respectively. Let $\{\sigma, \sigma'\}$ be an edge in $\text{FS}(X, Y)$. This means there

exists an edge $\{a, b\}$ in X such that $\{\sigma(a), \sigma(b)\}$ is an edge in Y , $\sigma(a) = \sigma'(b)$, $\sigma(b) = \sigma'(a)$, and $\sigma(c) = \sigma'(c)$ for all $c \in [n] \setminus \{a, b\}$. Since $\{a, b\}$ is also an edge in \tilde{X} and $\{\sigma(a), \sigma(b)\}$ is also an edge in \tilde{Y} , it follows that $\{\sigma, \sigma'\}$ is also an edge in $\text{FS}(\tilde{X}, \tilde{Y})$. \square

For $\sigma \in \mathfrak{S}_n$, we let $\text{sgn}(\sigma)$ denote the *sign* of σ , which is 1 if σ is an even permutation (i.e., a product of an even number of transpositions) and -1 if σ is an odd permutation (i.e., a product of an odd number of transpositions).

Proposition 2.2. *If X and Y are graphs on n vertices, then $\text{FS}(X, Y)$ is bipartite.*

Proof. We may assume $V(X) = V(Y) = [n]$. Let $A_n = \{\sigma \in \mathfrak{S}_n : \text{sgn}(\sigma) = 1\}$ denote the alternating group on n letters. The pair $\{A_n, \mathfrak{S}_n \setminus A_n\}$ is a bipartition of $\text{FS}(X, Y)$. \square

Proposition 2.3. *Let X and Y be graphs with $V(X) = V(Y) = [n]$. If $\varphi: [n] \rightarrow [n]$ is an automorphism of X , then the map $\varphi^*: \mathfrak{S}_n \rightarrow \mathfrak{S}_n$ given by $\varphi^*(\sigma) = \sigma \circ \varphi$ is an automorphism of $\text{FS}(X, Y)$.*

Proof. Let $\sigma, \sigma' \in \mathfrak{S}_n$. The pair $\{\sigma, \sigma'\}$ is an edge in $\text{FS}(X, Y)$ if and only if there exist $i, j \in [n]$ such that $\sigma' = \sigma \circ (i j)$, $\{i, j\} \in E(X)$, and $\{\sigma(i), \sigma(j)\} \in E(Y)$. The condition $\sigma' = \sigma \circ (i j)$ is equivalent to the condition $\varphi^*(\sigma') = \varphi^*(\sigma) \circ (\varphi^{-1}(i) \varphi^{-1}(j))$. Because φ^{-1} is an automorphism of X , the condition $\{i, j\} \in E(X)$ is equivalent to the condition that $\{\varphi^{-1}(i), \varphi^{-1}(j)\} \in E(X)$. We have $\{\sigma(i), \sigma(j)\} \in E(Y)$ if and only if $\{\varphi^*(\sigma)(\varphi^{-1}(i)), \varphi^*(\sigma)(\varphi^{-1}(j))\} \in E(Y)$ since these two pairs are actually equal. Thus, $\{\sigma, \sigma'\} \in E(\text{FS}(X, Y))$ if and only if $\{\varphi^*(\sigma), \varphi^*(\sigma')\} \in E(\text{FS}(X, Y))$. \square

The following simple proposition shows that in order to understand the connected components of graphs of the form $\text{FS}(X, Y)$, it suffices to understand what happens when X and Y are connected. Recall from Section 1.1 the definition of the graph Cartesian product $G_1 \square \dots \square G_r$.

Proposition 2.4. *Let X and Y be graphs on n vertices, and let X_1, \dots, X_r be the connected components of X . For $i \in [r]$, let $n_i = |V(X_i)|$. Let $\mathcal{OP}_{n_1, \dots, n_r}(Y)$ denote the collection of ordered set partitions (V_1, \dots, V_r) of $V(Y)$ such that $|V_i| = n_i$ for all $i \in [r]$. Then*

$$\text{FS}(X, Y) \cong \bigoplus_{(V_1, \dots, V_r) \in \mathcal{OP}_{n_1, \dots, n_r}(Y)} (\text{FS}(X_1, Y|_{V_1}) \square \dots \square \text{FS}(X_r, Y|_{V_r})).$$


Proof. For each bijection $\sigma: X \rightarrow Y$ and index $i \in [r]$, let $V_i^\sigma = \sigma(X_i)$. Then the ordered set partition $\rho_\sigma = (V_1^\sigma, \dots, V_r^\sigma)$ is in $\mathcal{OP}_{n_1, \dots, n_r}(Y)$. For each $\rho \in \mathcal{OP}_{n_1, \dots, n_r}(Y)$, let G_ρ be the induced subgraph of $\text{FS}(X, Y)$ whose vertex set consists of the bijections $\sigma: X \rightarrow Y$ such that $\rho_\sigma = \rho$. If σ' is obtained from σ by performing an (X, Y) -friendly swap (meaning $\{\sigma, \sigma'\} \in \text{FS}(X, Y)$), then we must have $\rho_\sigma = \rho_{\sigma'}$. It follows that every connected component of $\text{FS}(X, Y)$ is contained in one of the induced subgraphs G_ρ , so $\text{FS}(X, Y) \cong \bigoplus_{\rho \in \mathcal{OP}_{n_1, \dots, n_r}(Y)} G_\rho$.

Now fix $\rho = (V_1, \dots, V_r) \in \mathcal{OP}_{n_1, \dots, n_r}(Y)$. For $\sigma \in V(G_\rho)$, let $\psi(\sigma) = (\tau_1, \dots, \tau_r)$, where $\tau_k = \sigma|_{X_k}: X_k \rightarrow Y|_{V_k}$. The map ψ is a bijection from $V(G_\rho)$ to $V(\text{FS}(X_1, Y|_{V_1})) \times \dots \times V(\text{FS}(X_r, Y|_{V_r}))$. We want to show that ψ is an isomorphism from G_ρ to

$$\text{FS}(X_1, Y|_{V_1}) \square \dots \square \text{FS}(X_r, Y|_{V_r}).$$

To do this, choose some edge $\{\sigma, \sigma'\}$ of G_ρ . We know that σ' is obtained from σ by performing an (X, Y) -friendly swap across $\{i, j\}$ for some edge $\{i, j\} \in E(X)$. Let X_m be the connected component of X containing the edge $\{i, j\}$. Then $\psi(\sigma')$ is obtained from $\psi(\sigma)$ by performing an (X_m, Y_m) -friendly swap across $\{i, j\}$ to the m -th coordinate of $\psi(\sigma)$ and leaving all other coordinates of $\psi(\sigma)$ unchanged. This means that $\{\psi(\sigma), \psi(\sigma')\}$ is an edge in $\text{FS}(X_1, Y|_{V_1}) \square \dots \square \text{FS}(X_r, Y|_{V_r})$. This argument is reversible, so $\{\psi(\sigma), \psi(\sigma')\}$ is an edge in $\text{FS}(X_1, Y|_{V_1}) \square \dots \square \text{FS}(X_r, Y|_{V_r})$ if and only if $\{\sigma, \sigma'\}$ is an edge in G_ρ . \square

Let us now state Wilson's main result about the connected components of $\text{FS}(\text{Star}_n, Y)$. Note that Wilson considers only the case in which Y is biconnected (i.e., Y is connected and does not have a cut vertex). He also excludes cycle graphs, although these graphs are easy to handle separately. Indeed, it is not difficult to show that the graph $\text{FS}(\text{Star}_n, \text{Cycle}_n)$ has $(n-2)!$ connected components, each of which contains exactly $n(n-1)$ vertices. The connected components of $\text{FS}(\text{Star}_n, \text{Cycle}_n)$ correspond to cyclic orderings of $[n-1]$: indeed, by repeatedly swapping the vertex labeled n with its clockwise neighbor, one can obtain any other labeling of Cycle_n that witnesses the same cyclic ordering of $[n-1]$. There are $(n-2)!$ cyclic orderings of $[n-1]$; once one is chosen, a vertex σ in the corresponding connected component is uniquely determined by specifying $\sigma(n)$ and $\sigma(n-1)$. One other specific graph that Wilson must consider separately is the graph



Theorem 2.5 ([23]). *Let Y be a biconnected graph on $n \geq 3$ vertices that is not isomorphic to θ_0 or Cycle_n . If Y is not bipartite, then $\text{FS}(\text{Star}_n, Y)$ is connected. If Y is bipartite, then $\text{FS}(\text{Star}_n, Y)$ has exactly 2 connected components, each of size $n!/2$. The graph $\text{FS}(\text{Star}_7, \theta_0)$ has exactly 6 connected components.*

We now show that when both X and Y are separable, the resulting $\text{FS}(X, Y)$ has a very large number of components; it is for this reason that Wilson [23] restricted his investigation of $\text{FS}(\text{Star}_n, Y)$ to the case where Y is biconnected. By Proposition 2.4, it suffices to consider the case where X and Y are connected but separable.

Proposition 2.6. *Let X and Y be connected graphs, each on $n \geq 3$ vertices. Suppose $x_0 \in X$ and $y_0 \in Y$ are cut vertices such that the connected components of $X|_{V(X) \setminus \{x_0\}}$ are X_1, \dots, X_r and the connected components of $Y|_{V(Y) \setminus \{y_0\}}$ are Y_1, \dots, Y_s . Let $\mathcal{M} = \mathcal{M}(X, Y, x_0, y_0)$ denote the set of $r \times s$ matrices with nonnegative integer entries in which the i -th row sums to $|V(X_i)|$ and the j -th column sums to $|V(Y_j)|$. Then the number of connected components of $\text{FS}(X, Y)$ is at least $|\mathcal{M}|$.*

For more information about \mathcal{M} (including size estimates), we direct the reader to [3, 11] and the references therein. Note that we always have $|\mathcal{M}| \geq 2$ (so that $\text{FS}(X, Y)$ is disconnected) since $r, s \geq 2$ by the definition of a cut vertex. We also remark that, as will be apparent in the proof, this lower bound on the number of connected components of $\text{FS}(X, Y)$ is an equality when the graphs $X|_{V(X_i) \cup \{x_0\}}$ and $Y|_{V(Y_j) \cup \{y_0\}}$ are all complete graphs.

Proof. Let $\Sigma_0 = \{\sigma \in V(\text{FS}(X, Y)) : \sigma(x_0) = y_0\}$. For each $\sigma \in \Sigma_0$, define the $r \times s$ incidence matrix $M(\sigma) = (m_{i,j}(\sigma)) \in \mathcal{M}$ via $m_{i,j}(\sigma) = |\sigma(V(X_i)) \cap V(Y_j)|$. It is immediate that every $M \in \mathcal{M}$ arises as $M(\sigma)$ for some σ . Therefore, it suffices to show that $M(\sigma) = M(\sigma')$ whenever $\sigma, \sigma' \in \Sigma_0$ are in the same connected component of $\text{FS}(X, Y)$. Fix some $\sigma \in \Sigma_0$. We will actually prove the following stronger assertion for every $\tau \in V(\text{FS}(X, Y))$ in the same connected component as σ :

- if $\tau \in V(\Sigma_0)$, then $M(\tau) = M(\sigma)$;
- if X_{i_0} is the connected component of X containing $\tau^{-1}(y_0)$, then

$$|\tau(V(X_{i_0}) \cup \{x_0\}) \cap V(Y_j)| = m_{i_0,j}(\sigma) \quad \text{and} \quad |\tau(V(X_i)) \cap V(Y_j)| = m_{i,j}(\sigma)$$

for every $i \neq i_0$.

Let us assume that these conditions hold for some τ in the same connected component as σ . We will show that the claim also holds for every vertex τ' that is adjacent to τ in $\text{FS}(X, Y)$, which will complete the proof. Let τ' be obtained from τ by the application of an (X, Y) -friendly swap across the edge $\{x_1, x_2\} \in E(X)$.

First, suppose that $\tau \in V(\Sigma_0)$ (i.e., $\tau(x_0) = y_0$). If $x_0 \notin \{x_1, x_2\}$, then $\{x_1, x_2\}$ lies within a single connected component X_i of X , so $\tau' \in V(\Sigma_0)$ and $M(\tau') = M(\tau)$. It is also straightforward to check that the claim holds if instead (without loss of generality) $x_1 = x_0$ and $x_2 \in V(X_{i_0})$.

Second, suppose that $\tau^{-1}(y_0) \in V(X_{i_0})$. Again, the result is straightforward if $x_0 \notin \{x_1, x_2\}$ or if $y_0 \notin \{\tau(x_1), \tau(x_2)\}$. The only remaining case is (without loss of generality) where $x_0 = x_1$ and $\tau(x_2) = y_0$ (where we know that $x_2 \in X_{i_0}$). Here, the desired condition is again easy to check. □

Part of Wilson’s Theorem 2.5 says that $\text{FS}(\text{Star}_n, Y)$ is disconnected whenever Y is bipartite. The following result generalizes this to all bipartite X and Y . As before, let $\text{sgn}(\sigma)$ be 1 if σ is an even permutation and -1 if σ is an odd permutation.

Proposition 2.7. *If X and Y are bipartite graphs, each on $n \geq 3$ vertices, then $\text{FS}(X, Y)$ is disconnected.*

Proof. We may assume $V(X) = V(Y) = [n]$. Let $\{A_X, B_X\}$ be a bipartition of X , and let $\{A_Y, B_Y\}$ be a bipartition of Y . For each $\sigma \in \mathfrak{S}_n$, let

$$p(\sigma) = |\sigma(A_X) \cap A_Y| + \frac{\text{sgn}(\sigma) + 1}{2}.$$

We claim that if σ and σ' are adjacent in $\text{FS}(X, Y)$, then $p(\sigma)$ and $p(\sigma')$ have the same parity. Indeed, suppose σ' is obtained from σ by performing an (X, Y) -friendly swap across $\{i, j\}$. This means that $\{i, j\} \in E(X)$ and $\{\sigma(i), \sigma(j)\} \in E(Y)$. The numbers i and j belong to different sets in the bipartition $\{A_X, B_X\}$, and the numbers $\sigma(i)$ and $\sigma(j)$ belong to different sets in the bipartition $\{A_Y, B_Y\}$. Since $\sigma' = \sigma \circ (ij)$, it follows that $|\sigma(A_X) \cap A_Y|$ and $|\sigma'(A_X) \cap A_Y|$ differ by exactly 1 and that $(\text{sgn}(\sigma) + 1)/2$ and $(\text{sgn}(\sigma') + 1)/2$ differ by exactly 1. This proves

the claim that $p(\sigma) \equiv p(\sigma') \pmod{2}$. It follows that $p(\tau) \equiv p(\tau') \pmod{2}$ whenever τ and τ' are in the same connected component of $\text{FS}(X, Y)$.

Now fix $\tau \in \mathfrak{S}_n$. Since $n \geq 3$, there exist elements a and b of $[n]$ that lie in the same set in the bipartition $\{A_X, B_X\}$. Let $\tau' = \tau \circ (ab)$. Then $p(\tau)$ and $p(\tau')$ differ by 1, so τ and τ' must lie in different connected components of $\text{FS}(X, Y)$. \square

Remark 2.8. Note that Star_n is bipartite with bipartition $\{[n-1], \{n\}\}$. Suppose that $n \geq 3$ and that Y is a biconnected graph on the vertex set $[n]$ that is not isomorphic to Cycle_n or θ_0 . If Y is bipartite with bipartition $\{A_Y, B_Y\}$, then it follows from Theorem 2.5 that $\text{FS}(\text{Star}_n, Y)$ has exactly 2 connected components. Appealing to the proof of Proposition 2.7, we find that the connected components of $\text{FS}(\text{Star}_n, Y)$ must be precisely

$$\{\sigma \in \mathfrak{S}_n : p(n) \equiv 0 \pmod{2}\} \quad \text{and} \quad \{\sigma \in \mathfrak{S}_n : p(n) \equiv 1 \pmod{2}\},$$

where

$$p(\sigma) = |\sigma([n-1]) \cap A_Y| + \frac{\text{sgn}(\sigma) + 1}{2}.$$

Now suppose X is a graph with vertex set $[n]$ such that Star_n is a proper subgraph of X . This means that there exists an edge $\{i, j\} \in E(X)$ with $n \notin \{i, j\}$. There exists $\sigma \in \mathfrak{S}_n$ such that $\sigma(i)$ and $\sigma(j)$ are adjacent in Y . The permutation $\sigma' = \sigma \circ (ij)$ is adjacent to σ in $\text{FS}(X, Y)$, and $p(\sigma) \not\equiv p(\sigma') \pmod{2}$. This means that σ and σ' are in different connected components of $\text{FS}(\text{Star}_n, Y)$, but are in the same connected component of $\text{FS}(X, Y)$. Since each of the 2 connected components of $\text{FS}(\text{Star}_n, Y)$ is contained in a connected component of $\text{FS}(X, Y)$, it follows that $\text{FS}(X, Y)$ is connected. This fact demonstrates a sense in which Wilson's result is sharp.

3. Paths

In this section, we describe the structure of $\text{FS}(\text{Path}_n, Y)$. We will see that this characterization is closely related to a known result about Coxeter elements of Coxeter systems.

In what follows, we let G be a graph with vertex set $[n]$. We obtain an *orientation* of G by choosing a direction for each of its edges. An orientation is *acyclic* if it does not contain a directed cycle. Let $\text{Acyc}(G)$ denote the set of acyclic orientations of G . For every $\alpha \in \text{Acyc}(G)$, we obtain a poset $([n], \leq_\alpha)$ by declaring that $i \leq_\alpha j$ if and only if the directed graph α contains a directed path starting at the vertex i and ending at the vertex j . (When $i = j$, we can use the 1-vertex path with no edges.) We write $\mathcal{L}(\alpha)$ for the set of linear extensions of $([n], \leq_\alpha)$, as defined in the introduction. For each permutation $\sigma \in \mathfrak{S}_n$, there is a unique acyclic orientation $\alpha_G(\sigma) \in \text{Acyc}(G)$ such that $\sigma \in \mathcal{L}(\alpha_G(\sigma))$. Indeed, $\alpha_G(\sigma)$ is obtained by directing each edge $\{i, j\}$ of G from i to j if and only if $\sigma^{-1}(i) < \sigma^{-1}(j)$.

Theorem 3.1. *Let Y be a graph with vertex set $[n]$. For each $\alpha \in \text{Acyc}(\overline{Y})$, choose a linear extension $\sigma_\alpha \in \mathcal{L}(\alpha)$. Let H_α be the connected component of $\text{FS}(\text{Path}_n, Y)$ containing σ_α . The connected component H_α depends only on α (not on the specific choice of σ_α), and its vertex set*




Figure 3.1: A graph G labeled according to the permutation $\sigma = 53142$. Each vertex i is labeled with a plain black label i and a circled red label $\sigma^{-1}(i)$. The edges are oriented so as to form the acyclic orientation $\alpha_G(\sigma)$.

is $\mathcal{L}(\alpha)$. Moreover,

$$\text{FS}(\text{Path}_n, Y) = \bigoplus_{\alpha \in \text{Acyc}(\bar{Y})} H_\alpha.$$

Proof. Choose distinct $\sigma, \sigma' \in \mathfrak{S}_n$, and let $\alpha = \alpha_{\bar{Y}}(\sigma)$. The permutations σ and σ' are adjacent in $\text{FS}(\text{Path}_n, Y)$ if and only if there exists $i \in [n - 1]$ such that $\{\sigma(i), \sigma(i + 1)\}$ is an edge in Y and $\sigma' = \sigma \circ (i \ i + 1)$. Notice that in the poset $([n], \leq_\alpha)$, the element $\sigma(i + 1)$ either covers $\sigma(i)$ or is incomparable to $\sigma(i)$. This means that $\sigma(i)$ and $\sigma(i + 1)$ are adjacent in Y (i.e., nonadjacent in \bar{Y}) if and only if they are incomparable in $([n], \leq_\alpha)$. Letting t_1, \dots, t_{n-1} be the toggle operators from Proposition 1.3 (with respect to the poset $([n], \leq_\alpha)$), we see that σ and σ' are adjacent if and only if $\sigma' = t_i(\sigma)$ for some $i \in [n - 1]$. It follows that if σ and σ' are adjacent in $\text{FS}(\text{Path}_n, Y)$, then $\alpha_{\bar{Y}}(\sigma) = \alpha_{\bar{Y}}(\sigma')$. Using Proposition 1.3, we see that σ and σ' are in the same connected component of $\text{FS}(\text{Path}_n, Y)$ if and only if $\alpha_{\bar{Y}}(\sigma) = \alpha_{\bar{Y}}(\sigma')$. This is equivalent to the statement of the theorem. \square

It is well known that the number of acyclic orientations of a graph G is equal to the evaluation $T_G(2, 0)$ of the Tutte polynomial of G (see [6] for the definition of the Tutte polynomial). Furthermore, a graph with at least 1 edge has at least 2 acyclic orientations. Hence, we have the following corollary.

Corollary 3.2. *Let Y be a graph with vertex set $[n]$. The number of connected components of $\text{FS}(\text{Path}_n, Y)$ is $T_{\bar{Y}}(2, 0)$. In particular, $\text{FS}(\text{Path}_n, Y)$ is connected if and only if $Y = K_n$.*

Recall that a *Coxeter system* is a pair (W, S) , where W is a group with generating set $S = \{s_1, \dots, s_n\}$ and presentation $W = \langle S : (s_i s_j)^{m_{i,j}} = 1 \rangle$. Here, the exponents $m_{i,j}$ are elements of $\{1, 2, 3, \dots\} \cup \{\infty\}$ such that $m_{ii} = 1$ for all $i \in [n]$ and $m_{i,j} = m_{j,i} \geq 2$ whenever $i \neq j$. Note that the elements s_i and s_j commute if and only if $m_{i,j} \leq 2$. The *Coxeter graph*¹ associated to the Coxeter system (W, S) is the simple graph with vertex set S in which vertices s_i and s_j are adjacent if and only if $m_{i,j} \geq 3$ (i.e., $s_i s_j \neq s_j s_i$). A *Coxeter element* (several authors say *standard Coxeter element*) of (W, S) is an element of W of the form $s_{\sigma(1)} \dots s_{\sigma(n)}$, where $\sigma \in \mathfrak{S}_n$.

¹Coxeter graphs often have edge labels that encode the exponents $m_{i,j}$, but we will ignore those labels here.

Now let Y be a graph with vertex set $[n]$. There exists a Coxeter system (W, S) whose Coxeter graph is \bar{Y} , where we identify the vertex $i \in [n] = V(\bar{Y})$ with the element $s_i \in S$. With this identification, every permutation $\sigma \in \mathfrak{S}_n$ gives rise to a word $s_{\sigma(1)} \dots s_{\sigma(n)}$, which represents a Coxeter element of (W, S) . Two such words represent the same element of W if and only if one can be obtained from the other by repeatedly applying the commutation relations $s_i s_j = s_j s_i$, which hold when i and j are adjacent in Y . Applying such a commutation relation to a word $s_{\sigma(1)} \dots s_{\sigma(n)}$ means that we swap the factors $s_{\sigma(i)}$ and $s_{\sigma(i+1)}$ for some $i \in [n-1]$ such that $\{\sigma(i), \sigma(i+1)\}$ is an edge in Y . This corresponds precisely to applying a (Path_n, Y) -friendly swap to the permutation σ . Hence, the Coxeter elements $s_{\sigma(1)} \dots s_{\sigma(n)}$ and $s_{\sigma'(1)} \dots s_{\sigma'(n)}$ are equal if and only if σ and σ' are in the same connected component of $\text{FS}(\text{Path}_n, Y)$. It follows that Theorem 3.1 is equivalent to the following standard theorem about Coxeter elements (see [7, 10]).

Theorem 3.3 ([7, 10]). *Let (W, S) be a Coxeter system with Coxeter graph G , and write $S = \{s_1, \dots, s_n\}$. Identify each vertex s_i of G with the element i of $[n]$. For each acyclic orientation $\alpha \in \text{Acyc}(G)$, choose a linear extension σ_α of $([n], \leq_\alpha)$. The Coxeter element $s_{\sigma_\alpha(1)} \dots s_{\sigma_\alpha(n)}$ depends only on α , not on the specific linear extension σ_α . Also, the map $\alpha \mapsto s_{\sigma_\alpha(1)} \dots s_{\sigma_\alpha(n)}$ is a bijection from $\text{Acyc}(G)$ to the set of Coxeter elements of (W, S) .*

It is worth mentioning that Theorem 3.3, and hence Theorem 3.1, can also be interpreted outside of a Coxeter-theoretic framework using certain Cartier–Foata traces [7].

We end this section with a brief description of how connected components of $\text{FS}(\text{Path}_n, Y)$ inherit their structure from $\text{FS}(\text{Path}_n, K_n)$. It is immediate that every edge in $\text{FS}(\text{Path}_n, Y)$ is also an edge in $\text{FS}(\text{Path}_n, K_n)$. On the other hand, suppose $\sigma, \tau \in \mathfrak{S}_n$ are in the same connected component of $\text{FS}(\text{Path}_n, Y)$ and that they are adjacent in $\text{FS}(\text{Path}_n, K_n)$. There must exist $i \in [n-1]$ such that $\tau = \sigma \circ (i \ i+1)$. Since σ and τ are in the same connected component of $\text{FS}(\text{Path}_n, Y)$, Theorem 3.1 tells us that $\alpha_{\bar{Y}}(\sigma) = \alpha_{\bar{Y}}(\tau)$. This implies that $\sigma(i)$ and $\sigma(i+1)$ are adjacent in Y , so τ is obtained from σ by performing a (Path_n, Y) -friendly swap across $\{i, i+1\}$. Hence, σ and τ are adjacent in $\text{FS}(\text{Path}_n, Y)$. This yields the following corollary to Theorem 3.1.

Corollary 3.4. *Let Y be a graph with vertex set $[n]$. Let H_α be as in Theorem 3.1, so that*

$$\text{FS}(\text{Path}_n, Y) = \bigoplus_{\alpha \in \text{Acyc}(\bar{Y})} H_\alpha.$$

Then $H_\alpha = \text{FS}(\text{Path}_n, K_n)|_{\mathcal{L}(\alpha)}$ for every $\alpha \in \text{Acyc}(\bar{Y})$.

Remark 3.5. Note that $\text{FS}(\text{Path}_n, K_n)$ is the Hasse diagram of the weak order on \mathfrak{S}_n . A subset A of a poset P is called *convex* if for all $a, c \in A$ and all $b \in P$ satisfying $a <_P b <_P c$, we have $b \in A$. Björner and Wachs [5] showed that a subset $A \subseteq \mathfrak{S}_n$ is a convex subset of the weak order on \mathfrak{S}_n if and only if A is the set of linear extensions of a poset. It follows from Theorem 3.1 that the set of vertices of a connected component of $\text{FS}(\text{Path}_n, Y)$ is a convex subset of the weak order on \mathfrak{S}_n .

4. Cycles

In this section, we describe the structure of $\text{FS}(\text{Cycle}_n, Y)$. We will see that many of the results in this section are similar in form to those in the previous section, with acyclic orientations and posets replaced by their appropriate toric analogues. One of the key components of the proof of Theorem 3.1 was the fundamental Proposition 1.3, which states that the group of bijections generated by the toggles t_1, \dots, t_{n-1} acts transitively on the set of linear extensions of an n -element poset. Some new ideas will be needed to establish the theorems in this section.

Let G be a graph with vertex set $[n]$. A *source* of an acyclic orientation α of G is a vertex of in-degree 0; a *sink* of α is a vertex of out-degree 0. If v is a source or a sink of α , then we can obtain a new acyclic orientation of G by reversing the directions of all edges incident to v . We call this operation a *flip*. Two acyclic orientations $\alpha, \alpha' \in \text{Acyc}(G)$ are *torically equivalent*, denoted $\alpha \sim \alpha'$, if α' can be obtained from α via a sequence of flips. The equivalence classes in $\text{Acyc}(G)/\sim$ are called *toric acyclic orientations*. We denote the toric acyclic orientation containing the acyclic orientation α by $[\alpha]_\sim$.

Toric acyclic orientations have been studied in many different forms [8, 12, 16, 18, 20, 21]; the article [10] formalizes a systematic framework for their investigation. One of the reasons for the use of the word “toric” stems from their connection with hyperplane arrangements. Indeed, consider the *graphical arrangement* of the graph G , which is the hyperplane arrangement $\mathcal{A}(G)$ in \mathbb{R}^n consisting of the hyperplanes of the form $\{x \in \mathbb{R}^n : x_i = x_j\}$ for all edges $\{i, j\}$ of G . Let $\mathcal{A}_{\text{tor}}(G) = \pi(\mathcal{A}(G))$, where $\pi: \mathbb{R}^n \rightarrow \mathbb{R}^n/\mathbb{Z}^n$ is the natural projection map. The *toric chambers* of $\mathcal{A}_{\text{tor}}(G)$ are the connected components of $(\mathbb{R}^n/\mathbb{Z}^n) \setminus \mathcal{A}_{\text{tor}}(G)$. As shown in [10], there is a natural one-to-one correspondence between the toric chambers of $\mathcal{A}_{\text{tor}}(G)$ and the toric acyclic orientations of G . This is analogous to the standard fact that the chambers of $\mathcal{A}(G)$ (the connected components of $\mathbb{R}^n \setminus \mathcal{A}(G)$) are in bijection with $\text{Acyc}(G)$.

Another (related) motivation for the term “toric” comes from observing that flips encode what happens to the acyclic orientation associated to a permutation when we cyclically shift the permutation. To make this more precise, we let $\varphi: [n] \rightarrow [n]$ be the cyclic permutation given by $\varphi(i) = i + 1 \pmod n$ and consider the map $\varphi^*: \mathfrak{S}_n \rightarrow \mathfrak{S}_n$ defined by $\varphi^*(\sigma) = \sigma \circ \varphi$. Recall that $\alpha_G(\sigma)$ denotes the unique acyclic orientation of G such that $\sigma \in \mathcal{L}(\alpha_G(\sigma))$. The vertex $\sigma(1)$ of G is a source of $\alpha_G(\sigma)$. It is not hard to show that the acyclic orientation $\alpha_G(\varphi^*(\sigma))$ is obtained from $\alpha_G(\sigma)$ by flipping the vertex $\sigma(1)$ from a source into a sink. Consequently, the acyclic orientations $\alpha_G((\varphi^*)^k(\sigma))$ for $0 \leq k \leq n-1$ are all torically equivalent. It is also helpful to keep in mind that the map φ^* has order n and that the acyclic orientation $\alpha_G((\varphi^*)^{-1}(\sigma)) = \alpha_G((\varphi^*)^{n-1}(\sigma))$ is obtained from $\alpha_G(\sigma)$ by flipping the vertex $\sigma(n)$ from a sink into a source. Note that φ is an automorphism of Cycle_n . Proposition 2.3 tells us that for every graph Y with vertex set $[n]$, the map φ^* is an automorphism of $\text{FS}(\text{Cycle}_n, Y)$. Given an induced subgraph H of $\text{FS}(\text{Cycle}_n, Y)$, we write $\varphi^*(H)$ for the induced subgraph of $\text{FS}(\text{Cycle}_n, Y)$ on the vertex set $\varphi^*(V(H))$.

We define a *linear extension*² of $[\alpha]_\sim$ to be a permutation σ such that there exists an acyclic

²Note that our notion of a linear extension of a toric acyclic orientation differs from the definition of a “toric total order” in [10]. Indeed, that article defines a toric total order of $[\alpha]_\sim$ to be a cyclic equivalence class $\{(\varphi^*)^k(\sigma) : 0 \leq k \leq n-1\}$ such that σ is (using our definition) a linear extension of $[\alpha]_\sim$.

orientation $\hat{\alpha} \in [\alpha]_{\sim}$ with $\sigma \in \mathcal{L}(\hat{\alpha})$. Letting $\mathcal{L}([\alpha]_{\sim})$ denote the set of linear extensions of $[\alpha]_{\sim}$, we have

$$\mathcal{L}([\alpha]_{\sim}) = \bigcup_{\hat{\alpha} \in [\alpha]_{\sim}} \mathcal{L}(\hat{\alpha}).$$

For every permutation σ , the unique toric acyclic orientation of G that has σ as a linear extension is $[\alpha_G(\sigma)]_{\sim}$.

We will also need a new equivalence relation on acyclic orientations of a graph. Suppose $\alpha \in \text{Acyc}(G)$ has a source u and a sink v such that u and v are not adjacent. We can simultaneously flip u into a sink and flip v into a source; we call this a *double flip*. We say two acyclic orientations $\alpha, \alpha' \in \text{Acyc}(G)$ are *double-flip equivalent*, denoted $\alpha \approx \alpha'$, if α' can be obtained from α via a sequence of double flips. Let $[\alpha]_{\approx}$ denote the equivalence class in $\text{Acyc}(G)/\approx$ that contains α . Note that every equivalence class in Acyc/\approx is a union of equivalence classes in Acyc/\sim because each double flip can be viewed as two individual flips performed one after the other; here, it is crucial that we require u and v to non-adjacent in the definition of a double flip. A *linear extension* of a double-flip equivalence class $[\alpha]_{\approx}$ is a permutation σ such that $\sigma \in \mathcal{L}(\hat{\alpha})$ for some $\hat{\alpha} \in [\alpha]_{\approx}$. Letting $\mathcal{L}([\alpha]_{\approx})$ denote the set of linear extensions of $[\alpha]_{\approx}$, we have

$$\mathcal{L}([\alpha]_{\approx}) = \bigcup_{\hat{\alpha} \in [\alpha]_{\approx}} \mathcal{L}(\hat{\alpha}).$$




Figure 4.1: Each red box encompasses a toric acyclic orientation of K_3 . Each blue box encompasses a double-flip equivalence class of K_3 . Note that each double-flip equivalence class contains a single acyclic orientation; it is impossible to perform a double flip on an acyclic orientation of a complete graph because sources and sinks are all adjacent. We have colored sinks purple and sources green.

The next theorem tells us that the connected components of $\text{FS}(\text{Cycle}_n, Y)$ are the sets of linear extensions of double-flip equivalence classes of the acyclic orientations of \bar{Y} . The basic idea behind this result is as follows. We already know by Theorem 3.1 that the connected components of $\text{FS}(\text{Path}_n, Y)$ are the sets of linear extensions of the acyclic orientations of \bar{Y} . Performing a (Cycle_n, Y) -friendly swap to a bijection σ across the edge $\{n, 1\}$ corresponds precisely to changing the source $\sigma(1)$ into a sink and changing the sink $\sigma(n)$ into a source. When we perform this friendly swap, the vertices $\sigma(1)$ and $\sigma(n)$ must be adjacent in Y , meaning they are not adjacent in \bar{Y} . This is precisely why we forbid u and v from being adjacent in the definition of a double flip. One could also view each double-flip equivalence class as a cyclic version of a Cartier–Foata trace [7], where commutations are additionally allowed to wrap around the end of the word.




Figure 4.2: This graph has only 1 toric acyclic orientation, which is the union of the 5 double-flip equivalence classes encompassed in the blue boxes. We have colored sinks purple and sources green.

Theorem 4.1. *Let Y be a graph with vertex set $[n]$. For each $[\alpha]_{\approx} \in \text{Acyc}(\bar{Y})/\approx$, choose a linear extension $\sigma_{[\alpha]_{\approx}} \in \mathcal{L}([\alpha]_{\approx})$. Let $H_{[\alpha]_{\approx}}$ be the connected component of $\text{FS}(\text{Cycle}_n, Y)$ containing $\sigma_{[\alpha]_{\approx}}$. The connected component $H_{[\alpha]_{\approx}}$ depends only on $[\alpha]_{\approx}$ (not on the specific choice of $\sigma_{[\alpha]_{\approx}}$), and its vertex set is $\mathcal{L}([\alpha]_{\approx})$. Moreover,*

$$\text{FS}(\text{Cycle}_n, Y) = \bigoplus_{[\alpha]_{\approx} \in \text{Acyc}(\bar{Y})/\approx} H_{[\alpha]_{\approx}}.$$

Proof. The theorem is equivalent to the statement that two permutations $\sigma, \sigma' \in \mathfrak{S}_n$ are in the same connected component of $\text{FS}(\text{Cycle}_n, Y)$ if and only if $\alpha_{\bar{Y}}(\sigma) \approx \alpha_{\bar{Y}}(\sigma')$. First, suppose σ and σ' are adjacent in $\text{FS}(\text{Cycle}_n, Y)$. This means that σ' is obtained from σ by performing a (Cycle_n, Y) -friendly swap across an edge $\{i, j\} \in E(\text{Cycle}_n)$. If $i \leq n - 1$ and $j = i + 1$, then this (Cycle_n, Y) -friendly swap is also a (Path_n, Y) -friendly swap, so σ and σ' are adjacent in $\text{FS}(\text{Path}_n, Y)$. By Theorem 3.1, this implies that $\alpha_{\bar{Y}}(\sigma) = \alpha_{\bar{Y}}(\sigma')$. Now suppose $\{i, j\} = \{n, 1\}$. The vertices $\sigma(1)$ and $\sigma(n)$ are adjacent in Y , so they are not adjacent in \bar{Y} . The former vertex is a source in $\alpha_{\bar{Y}}(\sigma)$, and the latter is a sink in $\alpha_{\bar{Y}}(\sigma)$. It is straightforward to see that $\alpha_{\bar{Y}}(\sigma')$ is obtained from $\alpha_{\bar{Y}}(\sigma)$ by performing a double flip that simultaneously flips $\sigma(1)$ from a source to a sink and flips $\sigma(n)$ from a sink to a source. Consequently, $\alpha_{\bar{Y}}(\sigma) \approx \alpha_{\bar{Y}}(\sigma')$. This shows that $\alpha_{\bar{Y}}(\sigma) \approx \alpha_{\bar{Y}}(\sigma')$ whenever σ and σ' are adjacent in $\text{FS}(\text{Cycle}_n, Y)$, which implies that $\alpha_{\bar{Y}}(\sigma) \approx \alpha_{\bar{Y}}(\sigma')$ whenever σ and σ' are in the same connected component of $\text{FS}(\text{Cycle}_n, Y)$.

To prove the converse, notice that the vertex set of a connected component of $\text{FS}(\text{Cycle}_n, Y)$ is a union of vertex sets of connected components of $\text{FS}(\text{Path}_n, Y)$. By Theorem 3.1, the vertex set of each connected component of $\text{FS}(\text{Path}_n, Y)$ is the set of linear extensions of some acyclic orientation of \bar{Y} . Therefore, we need only show that if $\alpha, \alpha' \in \text{Acyc}(\bar{Y})$ satisfy $\alpha \approx \alpha'$, then there exist $\tau \in \mathcal{L}(\alpha)$ and $\tau' \in \mathcal{L}(\alpha')$ such that τ and τ' are in the same connected component of $\text{FS}(\text{Cycle}_n, Y)$. To prove this, it suffices to consider the case in which α' is obtained from α by performing a double flip. Thus, let us assume that there exist a source u of α and a sink v of α (with $\{u, v\} \notin E(\bar{Y})$) such that α' is obtained from α by flipping u into a sink and flipping v into a source. It is straightforward to see that there is a linear extension τ of α such that $\tau(1) = u$

and $\tau(n) = v$. The vertices u and v are adjacent in Y (since they are not adjacent in \bar{Y}), so the operation that changes τ into $\tau' := \tau \circ (1\ n)$ is a (Cycle_n, Y) -friendly swap. Observe that τ' is a linear extension of α' . Since τ and τ' are adjacent in $\text{FS}(\text{Cycle}_n, Y)$, they are certainly in the same connected component. \square

Corollary 4.2. *Let G be a graph on n vertices. Given $\alpha \in \text{Acyc}(G)$, let $\vec{\alpha}$ be an acyclic orientation obtained from α by flipping a source into a sink. The equivalence class $[\vec{\alpha}]_{\approx}$ depends on only the equivalence class $[\alpha]_{\approx}$, not the specific representative α or the source in α that is flipped to obtain $\vec{\alpha}$. Therefore, the map $\Phi: \text{Acyc}(G)/\approx \rightarrow \text{Acyc}(G)/\approx$ defined by $\Phi([\alpha]_{\approx}) = [\vec{\alpha}]_{\approx}$ is well-defined.*

Proof. Let Y be a graph with vertex set $[n]$ such that \bar{Y} is isomorphic to G . Choose $\alpha, \alpha' \in \text{Acyc}(\bar{Y})$ such that $\alpha \approx \alpha'$. Let u and u' be sources of α and α' , respectively. Let $\vec{\alpha}$ (respectively, $\vec{\alpha}'$) be the acyclic orientation obtained from α (respectively, α') by flipping u (respectively, u') into a sink. There exist $\sigma \in \mathcal{L}(\alpha)$ and $\sigma' \in \mathcal{L}(\alpha')$ such that $\sigma(1) = u$ and $\sigma'(1) = u'$. Because $\alpha \approx \alpha'$, it follows from Theorem 4.1 that σ and σ' are in the same connected component of $\text{FS}(\text{Cycle}_n, Y)$. Because φ^* is an automorphism of $\text{FS}(\text{Cycle}_n, Y)$, the permutations $\varphi^*(\sigma)$ and $\varphi^*(\sigma')$ are in the same connected component of $\text{FS}(\text{Cycle}_n, Y)$. We now check that $\varphi^*(\sigma) \in \mathcal{L}(\vec{\alpha})$ and $\varphi^*(\sigma') \in \mathcal{L}(\vec{\alpha}')$. Using Theorem 4.1 once again, we find that $\vec{\alpha} \approx \vec{\alpha}'$. This proves that the map Φ is well-defined. \square

Example 4.3. Figures 4.1 and 4.2 show toric acyclic orientations in red boxes and double-flip equivalence classes in blue boxes. In each case, the double-flip equivalence classes within each toric acyclic orientation are cyclically ordered from left to right. The map Φ sends each double-flip equivalence class to the next double-flip equivalence class in this cyclic order.

Theorem 4.1 describes the connected components of $\text{FS}(\text{Cycle}_n, Y)$, but it is possible to say even more. Namely, we will obtain a description of these connected components that relies on only the equivalence relation \sim , not the double-flip equivalence relation \approx . We first handle the case in which \bar{Y} is connected, where we will see that each toric acyclic orientation of \bar{Y} corresponds to a union of n pairwise isomorphic connected components of $\text{FS}(\text{Cycle}_n, Y)$. We will then use this result to understand the connected components of $\text{FS}(\text{Cycle}_n, Y)$ when \bar{Y} is not necessarily connected.

Proposition 4.4. *Let Y be a graph on the vertex set $[n]$ such that \bar{Y} is connected. For each toric acyclic orientation $[\alpha]_{\sim} \in \text{Acyc}(\bar{Y})/\sim$, choose a linear extension $\sigma_{[\alpha]_{\sim}}$ of $[\alpha]_{\sim}$, and let $J_{[\alpha]_{\sim}}$ be the connected component of $\text{FS}(\text{Cycle}_n, Y)$ containing $\sigma_{[\alpha]_{\sim}}$. The graphs*

$$J_{[\alpha]_{\sim}}, \varphi^*(J_{[\alpha]_{\sim}}), \dots, (\varphi^*)^{n-1}(J_{[\alpha]_{\sim}})$$

are distinct, pairwise isomorphic connected components of $\text{FS}(\text{Cycle}_n, Y)$. Moreover,

$$\text{FS}(\text{Cycle}_n, Y) = \bigoplus_{[\alpha]_{\sim} \in \text{Acyc}(\bar{Y})/\sim} \bigoplus_{k=0}^{n-1} (\varphi^*)^k(J_{[\alpha]_{\sim}}).$$

Proof. The claim that the graphs $J_{[\alpha]_{\sim}}, \varphi^*(J_{[\alpha]_{\sim}}), \dots, (\varphi^*)^{n-1}(J_{[\alpha]_{\sim}})$ are pairwise isomorphic connected components of $\text{FS}(\text{Cycle}_n, Y)$ follows immediately from the fact that φ^* is an automorphism of $\text{FS}(\text{Cycle}_n, Y)$ (by Proposition 2.3). Each equivalence class in $\text{Acyc}(\overline{Y})/\sim$ is a union of equivalence classes in $\text{Acyc}(\overline{Y})/\approx$, so it follows from Theorem 4.1 that

$$\text{FS}(\text{Cycle}_n, Y) = \bigoplus_{[\alpha]_{\sim} \in \text{Acyc}(\overline{Y})/\sim} \text{FS}(\text{Cycle}_n, Y)|_{\mathcal{L}([\alpha]_{\sim})}.$$

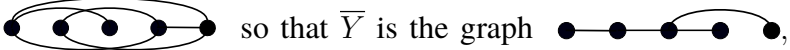
Let us fix a toric acyclic orientation $[\alpha]_{\sim} \in \text{Acyc}(\overline{Y})/\sim$ and a linear extension $\sigma_{[\alpha]_{\sim}} \in \mathcal{L}([\alpha]_{\sim})$. Let $U = \bigcup_{k=0}^{n-1} V((\varphi^*)^k(J_{[\alpha]_{\sim}}))$. We have seen that applying φ^* to a linear extension of an acyclic orientation $\beta \in \text{Acyc}(\overline{Y})$ produces a new permutation whose corresponding acyclic orientation is obtained from β by applying a flip; it follows that $(\varphi^*)^k(\sigma_{[\alpha]_{\sim}})$ is a linear extension of $[\alpha]_{\sim}$ for every $0 \leq k \leq n - 1$. In particular, $(\varphi^*)^k(\sigma_{[\alpha]_{\sim}}) \in V((\varphi^*)^k(J_{[\alpha]_{\sim}})) \cap \mathcal{L}([\alpha]_{\sim})$ for every $0 \leq k \leq n - 1$. Because $\text{FS}(\text{Cycle}_n, Y)|_{\mathcal{L}([\alpha]_{\sim})}$ is a union of connected components of $\text{FS}(\text{Cycle}_n, Y)$, we must have $U \subseteq \mathcal{L}([\alpha]_{\sim})$. We need to show that the reverse inclusion $\mathcal{L}([\alpha]_{\sim}) \subseteq U$ holds and that the subgraphs $J_{[\alpha]_{\sim}}, \varphi^*(J_{[\alpha]_{\sim}}), \dots, (\varphi^*)^{n-1}(J_{[\alpha]_{\sim}})$ are distinct.

Define a *move* to be an operation that changes a permutation τ into a permutation τ' such that either $\{\tau, \tau'\} \in E(\text{FS}(\text{Cycle}_n, Y))$ or $\tau' = \varphi^*(\tau)$. Say two permutations are *move-equivalent* if there is a sequence of moves transforming the first permutation into the second. Using the fact that $(\varphi^*)^{-1} = (\varphi^*)^{n-1}$, we see that move-equivalence is a genuine equivalence relation. Applying a move to an element of U produces another element of U . Therefore, in order to prove that $\mathcal{L}([\alpha]_{\sim}) \subseteq U$, it suffices to show that any two permutations in $\mathcal{L}([\alpha]_{\sim})$ are move-equivalent. If two permutations in $\mathcal{L}([\alpha]_{\sim})$ are linear extensions of the same acyclic orientation of \overline{Y} , then they are certainly move-equivalent because they are in the same connected component of $\text{FS}(\text{Cycle}_n, Y)$ by Theorem 4.1. Therefore, it suffices to show that for any two acyclic orientations $\alpha', \alpha'' \in [\alpha]_{\sim}$, there exist linear extensions $\sigma' \in \mathcal{L}(\alpha')$ and $\sigma'' \in \mathcal{L}(\alpha'')$ that are move-equivalent. In order to prove this, it suffices to prove it in the case where α'' is obtained from α' via a flip. Without loss of generality, we may assume α'' is obtained from α' by flipping a source u into a sink (otherwise, switch the roles of α' and α''). It is straightforward to see that there exists a linear extension σ' of α' such that $\sigma'(1) = u$. The permutation $\sigma'' = \varphi^*(\sigma')$ is a linear extension of α'' . By definition, σ' and σ'' are move-equivalent.

It remains to prove that $J_{[\alpha]_{\sim}}, \varphi^*(J_{[\alpha]_{\sim}}), \dots, (\varphi^*)^{n-1}(J_{[\alpha]_{\sim}})$ are distinct. This will follow if we can show that for every $\sigma \in \mathfrak{S}_n$ and every $k \in [n - 1]$, the vertices σ and $(\varphi^*)^k(\sigma)$ lie in different connected components of $\text{FS}(\text{Cycle}_n, Y)$. To prove this, it is helpful to imagine the cycle graph as lying in the plane with the vertices $1, \dots, n$ listed clockwise in this order. As in the introduction, we imagine n people labeled with the numbers $1, \dots, n$ standing on the vertices of the cycle graph, with the person labeled $\sigma(i)$ on the vertex i . Two people are friends with each other if and only if their labels are adjacent in Y , and two friends can swap places with each other whenever they are at adjacent positions in the cycle.

Now suppose by way of contradiction that there is a sequence of (Cycle_n, Y) -friendly swaps that changes the configuration of people given by σ to the configuration given by $(\varphi^*)^k(\sigma)$. The overall effect of this transformation is that each person moves k spaces counterclockwise. Imagine assigning a weight to each person so that a person's weight increases by 1 every time they move one space counterclockwise and decreases by 1 every time they move one space clockwise.

Assume each person's weight starts at 0. After the sequence of (Cycle_n, Y) -friendly swaps that has the effect of moving each person k spaces counterclockwise, each person's weight will be congruent to k modulo n . Since each swap increases one person's weight by 1 and decreases another person's weight by 1, the total weight of all people is always 0. Therefore, at least one person must have a positive weight in the end, and at least one person must have a negative weight in the end. Observe that no person has weight 0 at the end because $k \not\equiv 0 \pmod{n}$. Let A^+ be the (nonempty) set of people whose weights are positive in the end, and let A^- be the (nonempty) set of people whose weights are negative in the end. It is clear that if $p^+ \in A^+$ and $p^- \in A^-$, then at some point during the sequence of swaps, the person p^+ must have swapped places with p^- . This means that every person in A^+ is friends with every person in A^- . Using the labels to identify the people with vertices of Y , we see that every person in A^+ is adjacent in Y to every person in A^- . Since $A^+ \cup A^- = V(Y)$, this contradicts the assumption that \bar{Y} is connected. \square

Example 4.5. Suppose $Y = \text{---} \bullet \bullet \bullet \bullet \text{---}$ so that \bar{Y} is the graph , whose acyclic orientations are depicted in Figure 4.2. Identify the vertices of Y , read from left to right, with the numbers 1, 2, 3, 4, 5. Theorem 4.1 tells us that the connected components of $\text{FS}(\text{Cycle}_5, Y)$ correspond to the double-flip equivalence classes of \bar{Y} (the blue boxes in Figure 4.2). There is only 1 toric acyclic orientation $[\alpha]_{\sim}$ in $\text{Acyc}(\bar{Y})/\sim$, so we can choose $\sigma_{[\alpha]_{\sim}} = 12345$. In this case, $J_{[\alpha]_{\sim}}$ is the connected component of $\text{FS}(\text{Cycle}_5, Y)$ corresponding to the leftmost double-flip equivalence class in Figure 4.2. The set of vertices in $J_{[\alpha]_{\sim}}$ is the set of linear extensions of that double-flip equivalence class, which is

$$\{12354, 12345, 52341, 25341, 52314, 25314, 52134, 25134, 21534, 54312, 45312, 41532, 54132, 14532, 45132, 51432, 15432, 24351, 42315, 24315, 42135, 24135, 21435, 42351\} \quad (4.1)$$

The second double-flip equivalence class from the left, which is obtained by applying the map Φ from Corollary 4.2 to the leftmost double-flip equivalence class, also has 24 linear extensions; these are precisely the permutations obtained by applying φ^* to the permutations in (4.1). In general, applying Φ to a double-flip equivalence class $[\beta]_{\approx}$ yields the next double-flip equivalence class after $[\beta]_{\approx}$ in the left-to-right cyclic order; the linear extensions of $\Phi([\beta]_{\approx})$ are obtained by applying φ^* to the linear extensions of $[\beta]_{\approx}$.

The next corollary tells us about an interesting relationship between the sets $\text{Acyc}(G)/\sim$ and $\text{Acyc}(G)/\approx$ for an arbitrary graph G . It is worth noting that the proof of this fact, which is not obvious *a priori*, passes through the analysis of the graph $\text{FS}(\text{Cycle}_n, Y)$ with $Y = \bar{G}$.

Corollary 4.6. *Let G be a connected graph on n vertices, and let $\Phi: \text{Acyc}(G)/\approx \rightarrow \text{Acyc}(G)/\approx$ be the map from Corollary 4.2. For each $\alpha \in \text{Acyc}(G)$, the equivalence class $[\alpha]_{\sim}$ is the disjoint union*

$$[\alpha]_{\sim} = \bigsqcup_{k=0}^{n-1} \Phi^k([\alpha]_{\approx}).$$

Proof. Let $Y = \bar{G}$. We saw in the proof of Corollary 4.2 that applying Φ to an equivalence class $[\beta]_{\approx}$ corresponds to applying φ^* to the linear extensions of β . Preserving the notation

from Proposition 4.4, we can choose $\sigma_{[\alpha]_{\sim}}$ to be a linear extension of α so that $[\alpha]_{\sim}$ is the set of vertices of $J_{[\alpha]_{\sim}}$. Therefore, the desired result is equivalent to the part of Proposition 4.4 stating that

$$\text{FS}(\text{Cycle}_n, Y)|_{\mathcal{L}([\alpha]_{\sim})} = \bigoplus_{k=0}^{n-1} (\varphi^*)^k(J_{[\alpha]_{\sim}}). \quad \square$$

We can now proceed to describe the connected components of $\text{FS}(\text{Cycle}_n, Y)$ without assuming that \bar{Y} is connected. The reader may find it helpful to refer to Example 4.8 and Figure 4.3 while reading the following proof.

Theorem 4.7. *Let Y be a graph on the vertex set $[n]$. Let n_1, \dots, n_r denote the sizes of the connected components of \bar{Y} , and let $\nu = \text{gcd}(n_1, \dots, n_r)$. For each toric acyclic orientation $[\alpha]_{\sim} \in \text{Acyc}(\bar{Y})/\sim$, choose a linear extension $\sigma_{[\alpha]_{\sim}}$ of $[\alpha]_{\sim}$, and let $J_{[\alpha]_{\sim}}$ be the connected component of $\text{FS}(\text{Cycle}_n, Y)$ containing $\sigma_{[\alpha]_{\sim}}$. The graphs $J_{[\alpha]_{\sim}}, \varphi^*(J_{[\alpha]_{\sim}}), \dots, (\varphi^*)^{\nu-1}(J_{[\alpha]_{\sim}})$ are distinct, pairwise isomorphic connected components of $\text{FS}(\text{Cycle}_n, Y)$. Moreover,*

$$\text{FS}(\text{Cycle}_n, Y) = \bigoplus_{[\alpha]_{\sim} \in \text{Acyc}(\bar{Y})/\sim} \bigoplus_{k=0}^{\nu-1} (\varphi^*)^k(J_{[\alpha]_{\sim}}).$$

Proof. As in the proof of Proposition 4.4, the assertions that $J_{[\alpha]_{\sim}}, \varphi^*(J_{[\alpha]_{\sim}}), \dots, (\varphi^*)^{\nu-1}(J_{[\alpha]_{\sim}})$ are pairwise isomorphic connected components of $\text{FS}(\text{Cycle}_n, Y)$ and that

$$\text{FS}(\text{Cycle}_n, Y) = \bigoplus_{[\alpha]_{\sim} \in \text{Acyc}(\bar{Y})/\sim} \text{FS}(\text{Cycle}_n, Y)|_{\mathcal{L}([\alpha]_{\sim})}$$

are immediate. Let us fix a toric acyclic orientation $[\alpha]_{\sim} \in \text{Acyc}(\bar{Y})/\sim$ and a linear extension $\sigma_{[\alpha]_{\sim}} \in \mathcal{L}([\alpha]_{\sim})$. We may assume that $\sigma_{[\alpha]_{\sim}}$ is actually a linear extension of the specific acyclic orientation α . The same argument used in the proof of Proposition 4.4 shows that the vertex sets of $J_{[\alpha]_{\sim}}, \varphi^*(J_{[\alpha]_{\sim}}), \dots, (\varphi^*)^{\nu-1}(J_{[\alpha]_{\sim}})$ are all contained in $\mathcal{L}([\alpha]_{\sim})$. We need to show that the subgraphs $J_{[\alpha]_{\sim}}, \varphi^*(J_{[\alpha]_{\sim}}), \dots, (\varphi^*)^{\nu-1}(J_{[\alpha]_{\sim}})$ are distinct and that every element of $\mathcal{L}([\alpha]_{\sim})$ is a vertex of one of these ν subgraphs. Since $J_{[\alpha]_{\sim}}$ is the connected component of $\text{FS}(\text{Cycle}_n, Y)$ containing $\sigma_{[\alpha]_{\sim}}$ (which is a linear extension of α), it follows from Theorem 4.1 that $V(J_{[\alpha]_{\sim}}) = \mathcal{L}([\alpha]_{\sim})$. The proof of Corollary 4.2 tells us that the vertex set of $(\varphi^*)^k(J_{[\alpha]_{\sim}})$ is the set of linear extensions of $\Phi^k([\alpha]_{\sim})$. Therefore, we will be done if we can prove that $[\alpha]_{\sim}$ decomposes into the disjoint union

$$[\alpha]_{\sim} = \bigsqcup_{k=0}^{\nu-1} \Phi^k([\alpha]_{\sim}). \tag{4.2}$$

Let Z_1, \dots, Z_r denote the connected components of \bar{Y} , where each Z_i has size n_i . Let

$$\Phi_i: \text{Acyc}(Z_i)/\approx \rightarrow \text{Acyc}(Z_i)/\approx$$

be the map that flips a source of an acyclic orientation of Z_i into a sink, which is well-defined by Corollary 4.2. We now define yet another equivalence relation on $\text{Acyc}(\bar{Y})$. Say that a double

1 connected component; according to Theorem 4.1, this implies that there is only 1 double-flip equivalence class of acyclic orientations of \bar{Y} . Indeed, this is because Γ is generated by the element $\gamma_1 - \gamma_2 = (1, -1)$. We can reach any local-double-flip equivalence class from any other by repeatedly applying $\gamma_1 - \gamma_2$, and each application of $\gamma_1 - \gamma_2$ corresponds to applying a non-local double flip.




Figure 4.3: The action of Φ_1 and Φ_2 on local-double-flip equivalence classes.

Remark 4.9. Suppose G is a graph with connected components of sizes n_1, \dots, n_r , and let $\nu = \gcd(n_1, \dots, n_r)$. Letting $Y = \bar{G}$, we see immediately from Theorem 4.1 and Theorem 4.7 that each toric acyclic orientation of G is a disjoint union of exactly ν double-flip equivalence classes. *A priori*, it is not obvious that this should be the case.

Theorem 4.7 allows us to enumerate the connected components of $\text{FS}(\text{Cycle}_n, Y)$ in terms of the Tutte polynomial $T_{\bar{Y}}(x, y)$ of \bar{Y} ; this is analogous to Corollary 3.2.

Corollary 4.10. *Let Y be a graph with $n \geq 3$ vertices. Let Z_1, \dots, Z_r be the connected components of \bar{Y} , and let $\nu = \gcd(|V(Z_1)|, \dots, |V(Z_r)|)$. Then the number of connected components of $\text{FS}(\text{Cycle}_n, Y)$ is $T_{\bar{Y}}(1, 0)\nu$.*

Proof. It is known [10] that $T_{\bar{Y}}(1, 0)$ is the number of toric acyclic orientations of \bar{Y} , so the enumeration of the connected components follows from Theorem 4.7. □

Remark 4.11. If \bar{Y} is connected, then the number of connected components of $\text{FS}(\text{Cycle}_n, Y)$ is $T_{\bar{Y}}(1, 0)n$. It follows from a well-known theorem of Greene and Zaslavsky [14] that this is also the number of acyclic orientations of \bar{Y} with exactly 1 source.


For an even more concrete application of Theorem 4.7, we consider the special case in which \bar{Y} is a forest.

Corollary 4.12. *Let Y be a graph with $n \geq 3$ vertices such that \bar{Y} is a forest consisting of trees $\mathcal{T}_1, \dots, \mathcal{T}_r$, and let $\nu = \gcd(|V(\mathcal{T}_1)|, \dots, |V(\mathcal{T}_r)|)$. Then $\text{FS}(\text{Cycle}_n, Y)$ has ν connected components. These connected components are pairwise isomorphic, and they each contain $n!/\nu$ vertices. Moreover, each connected component of $\text{FS}(\text{Cycle}_n, Y)$ has an automorphism of order 2.*


Proof. We may assume $V(Y) = [n]$. It is known that $T_{\bar{Y}}(1, 0) = |\text{Acyc}(\bar{Y})/\sim| = 1$ because \bar{Y} is a forest. Consequently, the statements about the enumeration and sizes of the connected components of $\text{FS}(\text{Cycle}_n, Y)$ follow directly from Theorem 4.7.

Let H be one of the connected components of $\text{FS}(\text{Cycle}_n, Y)$. By Theorem 4.7, the connected components of $\text{FS}(\text{Cycle}_n, Y)$ are $H, \varphi^*(H), \dots, (\varphi^*)^{\nu-1}(H)$. We want to show that each connected component $(\varphi^*)^i(H)$ has an automorphism of order 2; since these connected components are pairwise isomorphic, it suffices to prove that H has an automorphism of order 2. The automorphism group of Cycle_n is the dihedral group of order $2n$; let ψ be one of the reflections in this group. Let $m \in \{0, \dots, \nu - 1\}$ be such that $\psi^*(H) = (\varphi^*)^m(H)$ (where, by abuse of notation, $\psi^*(H)$ denotes the connected component whose vertices form the set $\psi^*(V(H))$). Then $(\psi \circ \varphi^{-m})^*(H) = (\varphi^*)^{-m}(\psi^*(H)) = H$. Since $\psi \circ \varphi^{-m}$ is a reflection in the dihedral group of order $2n$, it has order 2. \square

Example 4.13. Let us revisit Example 1.2, which concerns the graph $\text{FS}(\text{Cycle}_5, Y)$ where



Notice that the complement of Y is the tree



Corollary 4.12 now explains why $\text{FS}(\text{Cycle}_5, Y)$ has 5 pairwise isomorphic connected components, each with an automorphism of order 2.

Corollary 4.14. *Let Y be a graph with $n \geq 3$ vertices. The graph $\text{FS}(\text{Cycle}_n, Y)$ is connected if and only if \bar{Y} is a forest consisting of trees $\mathcal{T}_1, \dots, \mathcal{T}_r$ such that $\gcd(|V(\mathcal{T}_1)|, \dots, |V(\mathcal{T}_r)|) = 1$.*

Proof. If \bar{Y} is a forest consisting of trees $\mathcal{T}_1, \dots, \mathcal{T}_r$, then Corollary 4.12 tells us that the graph $\text{FS}(\text{Cycle}_n, Y)$ is connected if and only if $\gcd(|V(\mathcal{T}_1)|, \dots, |V(\mathcal{T}_r)|) = 1$. Now suppose \bar{Y} is not a forest, meaning it contains a cycle of some length m with $3 \leq m \leq n$. Let Q be the graph consisting of $n - m$ isolated vertices and one cycle of length m . Since Q is a subgraph of \bar{Y} , the graph Y must be a subgraph of \bar{Q} . Our goal is to show that $\text{FS}(\text{Cycle}_n, Y)$ is disconnected; by Proposition 2.1, it suffices to prove that $\text{FS}(\text{Cycle}_n, \bar{Q})$ is disconnected. By Theorem 4.7, it suffices to show that Q has at least 2 toric acyclic orientations. As mentioned above, the number of toric acyclic orientations of Q is $T_Q(1, 0)$, where $T_Q(x, y)$ is the Tutte polynomial of Q . It is well known (and easy to prove from definitions) that $T_Q(x, y) = (x + x^2 + \dots + x^{m-1}) + y$, so $T_Q(1, 0) = m - 1 \geq 2$, as desired. \square

Recall that a graph on n vertices is called *Hamiltonian* if it has a subgraph isomorphic to Cycle_n . By combining Corollary 4.14 with Proposition 2.1, we obtain a sufficient condition to guarantee that $\text{FS}(X, Y)$ is connected whenever X is Hamiltonian: if X and Y are graphs on n vertices such that X is Hamiltonian and \bar{Y} is a forest consisting of trees $\mathcal{T}_1, \dots, \mathcal{T}_r$ with $\gcd(|V(\mathcal{T}_1)|, \dots, |V(\mathcal{T}_r)|) = 1$, then $\text{FS}(X, Y)$ is connected.

5. Sufficient Conditions for Connectivity

A *hereditary class* is a collection of (isomorphism types of) graphs that is closed under taking induced subgraphs. For example, for each fixed $d \geq 1$, the set of all graphs with maximum degree at most d is hereditary. Another example of a hereditary class is the set of all bipartite graphs. Recall that if G is a graph on n vertices, then a *Hamiltonian path* in G is a subgraph of G isomorphic to Path_n .

Our goal in this section is to find sufficient conditions on the graphs X and Y that guarantee $\text{FS}(X, Y)$ being connected. Our general setup will involve some fixed hereditary class \mathcal{H} . We will consider the case in which X contains a Hamiltonian path and \bar{Y} belongs to \mathcal{H} . We will state our main results in fairly general terms and then exhibit several specific applications. Our results rely heavily on Theorem 3.1, which tells us that the connected components of $\text{FS}(\text{Path}_n, Y)$ correspond to acyclic orientations of \bar{Y} . It will be convenient to assume that X has vertex set $[n]$ and that Path_n is a genuine subgraph of X (meaning $\{i, i + 1\} \in E(X)$ for all $i \in [n - 1]$). This does not sacrifice any generality because the isomorphism type of $\text{FS}(X, Y)$ depends only on the isomorphism types of X and Y . Proposition 2.1 tells us that the vertex set of each connected component of $\text{FS}(X, Y)$ is a union of vertex sets of connected components of $\text{FS}(\text{Path}_n, Y)$; Theorem 3.1 tells us that the vertex set of a connected component of $\text{FS}(\text{Path}_n, Y)$ is the set of linear extensions of the corresponding acyclic orientation of \bar{Y} .

In what follows, recall that if G is a graph with vertex set $[n]$ and $\sigma \in \mathfrak{S}_n$, then $\alpha_G(\sigma)$ denotes the unique acyclic orientation of G that has σ as a linear extension.

Theorem 5.1. *Let X and Y be graphs with $V(X) = V(Y) = [n]$, and suppose that Path_n is a subgraph of X . Suppose also that each connected component B of $\text{FS}(X, Y)$ contains some permutation σ_B such that n is a sink of $\alpha_{\bar{Y}}(\sigma_B)$. Then the number of connected components of $\text{FS}(X, Y)$ is at most the number of connected components of $\text{FS}(X|_{[n-1]}, Y|_{[n-1]})$.*

Proof. Let B_1, \dots, B_r be the connected components of $\text{FS}(X, Y)$, and let C_1, \dots, C_s be the connected components of $\text{FS}(X|_{[n-1]}, Y|_{[n-1]})$. Given $i \in [r]$, we know by hypothesis that there is a permutation σ_{B_i} in the connected component B_i such that n is a sink of $\alpha_{\bar{Y}}(\sigma_{B_i})$. Theorem 3.1 tells us that every linear extension of $\alpha_{\bar{Y}}(\sigma_{B_i})$ is in B_i . By replacing σ_{B_i} with a different linear extension of $\alpha_{\bar{Y}}(\sigma_{B_i})$ that sends n to n (such a linear extension certainly exists), we may assume without loss of generality that $\sigma_{B_i}(n) = n$. The acyclic orientation $\alpha_{\bar{Y}}(\sigma_{B_i})$ restricts to an acyclic orientation β_{B_i} of $\bar{Y}|_{[n-1]}$. Because Path_{n-1} is a subgraph of $X|_{[n-1]}$, Theorem 3.1 implies that all linear extensions of β_{B_i} belong to a single connected component $f(B_i)$ of $\text{FS}(X|_{[n-1]}, Y|_{[n-1]})$. We will show that the resulting map $f: \{B_1, \dots, B_r\} \rightarrow \{C_1, \dots, C_s\}$ is injective.

Suppose $f(B_i) = f(B_j) = C_k$. Then $\alpha_{\bar{Y}}(\sigma_{B_i})$ and $\alpha_{\bar{Y}}(\sigma_{B_j})$ both have n as a sink, and $\mathcal{L}(\beta_{B_i})$ and $\mathcal{L}(\beta_{B_j})$ are both subsets of $V(C_k)$. Because σ_{B_i} and σ_{B_j} send n to n , their restrictions $\sigma'_{B_i} = \sigma_{B_i}|_{[n-1]}$ and $\sigma'_{B_j} = \sigma_{B_j}|_{[n-1]}$ are linear extensions of β_{B_i} and β_{B_j} (respectively) and hence are both vertices in C_k . Thus, there is a sequence of $(X|_{[n-1]}, Y|_{[n-1]})$ -friendly swaps that transforms σ'_{B_i} into σ'_{B_j} . This sequence of $(X|_{[n-1]}, Y|_{[n-1]})$ -friendly swaps can be viewed as a sequence of (X, Y) -friendly swaps that transforms σ_{B_i} into σ_{B_j} , which implies that $B_i = B_j$, as desired. \square

We now investigate more carefully the conditions under which $\text{FS}(X, Y)$ is connected.

Theorem 5.2. *Let X and Y be graphs with $V(X) = V(Y) = [n]$, and suppose that Path_n is a subgraph of X . Suppose that Y is connected and that for every $(n-1)$ -vertex induced subgraph Y' of Y , the graph $\text{FS}(X|_{[n-1]}, Y')$ is connected. Then $\text{FS}(X, Y)$ is connected.*

Proof. The theorem will follow from Theorem 5.1 if we can show that each connected component of $\text{FS}(X, Y)$ contains a permutation such that the vertex n is a sink in the associated acyclic orientation of \bar{Y} . Fix a connected component B_i of $\text{FS}(X, Y)$; we say a vertex v of \bar{Y} is a \bar{Y} -sink relative to B_i if there exists a permutation in $V(B_i)$ whose associated acyclic orientation of \bar{Y} has v as a sink. Our goal is to prove that n is a \bar{Y} -sink relative to B_i ; we will actually prove the stronger fact that every vertex in \bar{Y} is a \bar{Y} -sink relative to B_i . First, note that at least one vertex of \bar{Y} must be a \bar{Y} -sink relative to B_i . Indeed, if $\sigma \in V(B_i)$, then $\sigma(n)$ is a \bar{Y} -sink relative to B_i . Because Y is connected, it now suffices to prove that if $\{y, z\} \in E(Y)$ and y is a \bar{Y} -sink relative to B_i , then z is also a \bar{Y} -sink relative to B_i .

Assume $\{y, z\} \in E(Y)$ and y is a \bar{Y} -sink relative to B_i . There is a permutation $\sigma \in V(B_i)$ such that y is a sink of $\alpha_{\bar{Y}}(\sigma)$. Theorem 3.1 tells us that every linear extension of $\alpha_{\bar{Y}}(\sigma)$ is a vertex in B_i , and one such linear extension must send n to y . Therefore, we may assume without loss of generality that $\sigma(n) = y$. Let Y' be the induced subgraph of Y on $[n] \setminus \{y\}$. The vertices in $\text{FS}(X|_{[n-1]}, Y')$ are the bijections from $[n-1]$ to $[n] \setminus \{y\}$, which are just the restrictions to $[n-1]$ of the bijections in \mathfrak{S}_n that send n to y . The hypothesis that $\text{FS}(X|_{[n-1]}, Y')$ is connected tells us that there is a sequence of $(X|_{[n-1]}, Y')$ -friendly swaps that transforms the vertex $\sigma' = \sigma|_{[n-1]}$ into a permutation τ' satisfying $\tau'(n-1) = z$. Let τ be the vertex in $\text{FS}(X, Y)$ that agrees with τ' on $[n-1]$ and satisfies $\tau(n) = y$. In particular, $\tau(n-1) = z$. Since $\{y, z\}$ is not an edge in \bar{Y} , it follows that z must be a sink in $\alpha_{\bar{Y}}(\tau)$. Finally, the same sequence of $(X|_{[n-1]}, Y')$ -friendly swaps that transforms σ' into τ' can be interpreted as a sequence of (X, Y) -friendly swaps that transforms σ into τ . This proves that $\tau \in V(B_i)$, so z is a \bar{Y} -sink relative to B_i , as desired. \square

It will be helpful to have a notion that captures the idea of extending a Hamiltonian path of a graph X and then adding additional edges. Thus, if X is a graph with a Hamiltonian path, then we define a *prolongation* of X to be a graph \tilde{X} such that:

- \tilde{X} contains a (not necessarily induced) subgraph $X^\#$ that is isomorphic to X ;
- \tilde{X} contains a Hamiltonian path that itself contains a Hamiltonian path of $X^\#$.

Theorem 5.2 has an immediate corollary in the language of hereditary classes and prolongations.

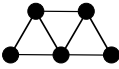
Corollary 5.3. *Let \mathcal{H} be a hereditary class. Let X be a graph on n_0 vertices with a Hamiltonian path, and suppose that $\text{FS}(X, Y)$ is connected for every $Y \in \mathcal{H}$ on n_0 vertices. If \tilde{X} is a prolongation of X with n vertices, then $\text{FS}(\tilde{X}, \tilde{Y})$ is connected for every $\tilde{Y} \in \mathcal{H}$ on n vertices.*

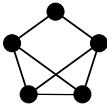
Proof. The proof is by induction on n . The case $n = n_0$ is obvious, so suppose $n > n_0$. By assumption, \tilde{X} contains a subgraph $X^\#$ that is isomorphic to X and a Hamiltonian path P that contains a Hamiltonian path of $X^\#$. Since $n > n_0$, we can find an endpoint p of the path P such that $\tilde{X}|_{V(\tilde{X}) \setminus \{p\}}$ is a prolongation of X on $n-1$ vertices. Let us identify $V(\tilde{X})$ with $[n]$ in such

a way that p is identified with n and P is identified with Path_n . Then $\tilde{X}|_{[n-1]}$ is a prolongation of X . Since \tilde{Y} is in the hereditary class \mathcal{H} , every $(n-1)$ -vertex induced subgraph of $\tilde{Y}|_{V(\tilde{Y}) \setminus \{y\}}$ is in \mathcal{H} . By our induction hypothesis, the graph $\text{FS}(X|_{[n-1]}, \tilde{Y}|_{V(\tilde{Y}) \setminus \{y\}})$ is connected for every $y \in V(\tilde{Y})$. If we identify $V(\tilde{Y})$ with $[n]$ in an arbitrary way, then it follows from Theorem 5.2 that $\text{FS}(\tilde{X}, \tilde{Y})$ is connected. \square


The set of all graphs Y such that \bar{Y} has maximum degree at most d forms a natural hereditary class. Equivalently, this hereditary class consists of all graphs Y with minimum degree at least $|V(Y)| - d - 1$. Given some $d \geq 1$ and a graph X on n_0 vertices, it requires only a finite search to determine whether or not $\text{FS}(X, Y)$ is connected for every graph Y on n_0 vertices with minimum degree at least $n_0 - d - 1$. Carrying out this finite computation for various d 's and X 's (with computer assistance) gives corollaries for infinite families of graphs. Each of these results would presumably require significant work to prove independently in an ad hoc way; Corollary 5.3 provides a unified framework for understanding them.

Corollary 5.4. *Let X be the triangle K_3 . For every prolongation \tilde{X} of X and for every graph \tilde{Y} on $|V(\tilde{X})|$ vertices with minimum degree at least $|V(\tilde{X})| - 2$, the graph $\text{FS}(\tilde{X}, \tilde{Y})$ is connected.*

Corollary 5.5. *Let X be the graph . For every prolongation \tilde{X} of X and for every graph \tilde{Y} on $|V(\tilde{X})|$ vertices with minimum degree at least $|V(\tilde{X})| - 3$, the graph $\text{FS}(\tilde{X}, \tilde{Y})$ is connected.*

Corollary 5.6. *Let X be the graph . For every prolongation \tilde{X} of X and for every graph \tilde{Y} on $|V(\tilde{X})|$ vertices with minimum degree at least $|V(\tilde{X})| - 3$, the graph $\text{FS}(\tilde{X}, \tilde{Y})$ is connected.*

Our motivation for investigating the phenomena of this section began with an attempt to understand $\text{FS}(X, Y)$ when X is the graph



More specifically, we are interested in characterizing the graphs Y such that $\text{FS}(X, Y)$ is connected. It follows from Corollary 5.4 that requiring Y to have minimum degree at least $n - 2$ is a sufficient condition to guarantee that $\text{FS}(X, Y)$ is connected. We will see in the next section that this condition is also necessary.

6. Necessary Conditions for Connectivity

We begin by describing a fairly general necessary condition for $\text{FS}(X, Y)$ to be connected in the case where X has a long path consisting of cut vertices and Y has a vertex of low degree. This result can be understood as an extension of Proposition 2.6.

Theorem 6.1. *Let X and Y be graphs on n vertices. Suppose $x_1 \dots x_d$ ($d \geq 1$) is a path in X , where x_1 and x_d are cut vertices and each of x_2, \dots, x_{d-1} has degree exactly 2. If the minimum degree of Y is at most d , then $\text{FS}(X, Y)$ is disconnected.*

Proof. Proposition 2.1 tells us that adding edges to Y cannot make $\text{FS}(X, Y)$ become disconnected, so we may assume that the minimum degree of Y is exactly d ; let $y_0 \in V(Y)$ be a vertex of degree d , and denote its neighbors by y_1, \dots, y_d . We now identify special vertices $x_0, x_{d+1} \in V(X)$ as follows. First, suppose $d = 1$. Since x_1 is a cut vertex of X , it has neighbors in multiple connected components of $X|_{V(X) \setminus \{x_1\}}$; let x_0 and x_2 be neighbors of x_1 in different connected components of $X|_{V(X) \setminus \{x_1\}}$. Second, suppose $d > 1$. Since x_1 is a cut vertex of X , it has a neighbor in a connected component of $X|_{V(X) \setminus \{x_1\}}$ that does not contain x_2 ; fix x_0 to be any such neighbor of x_1 . Similarly, since x_d is a cut vertex of X , it has a neighbor in a connected component of $X|_{V(X) \setminus \{x_d\}}$ that does not contain x_{d-1} ; fix x_{d+1} to be any such neighbor of x_d . Note that any path from x_0 to x_{d+1} must traverse $x_1 \dots x_d$. Note also that each of the vertices x_1, \dots, x_d is a cut vertex of X . Finally, let R denote the vertex set of the connected component of $X|_{V(X) \setminus \{x_1\}}$ that contains x_0 . See Figure 6.1.

Fix some bijection $\sigma \in \text{FS}(X, Y)$ such that $\sigma(x_0) = y_0$. Let B be the connected component of $\text{FS}(X, Y)$ that contains σ . We claim that every bijection τ in B satisfies $\tau(x_{d+1}) \neq y_0$; this will imply that $\text{FS}(X, Y)$ has multiple connected components. We will in fact establish the following stronger statement: If τ is in B , then either

- (i) $\tau^{-1}(y_0) \in R$; or
- (ii) $\tau^{-1}(y_0) = x_i$ for some $1 \leq i \leq d$, and at least i of the vertices y_1, \dots, y_d are contained in $\tau(R \cup \{x_1, \dots, x_{i-1}\})$.

Note that condition (i) holds for $\tau = \sigma$. Thus, it suffices to show that if this statement holds for some τ in B , then it also holds for any τ' that is obtained from τ by an (X, Y) -friendly swap. Suppose τ' is obtained from τ by an (X, Y) -friendly swap across the edge $\{u, v\} \in E(X)$ (so that $\{\tau(u), \tau(v)\} \in E(Y)$). We consider a number of case distinctions.

- (a) Suppose τ satisfies condition (i) and neither u nor v equals $\tau^{-1}(y_0)$. Then $(\tau')^{-1}(y_0) = \tau^{-1}(y_0) \in R$, so τ' satisfies condition (i).
- (b) Suppose τ satisfies condition (i) and $\tau^{-1}(y_0) \in \{u, v\}$. Without loss of generality, assume $\tau^{-1}(y_0) = u$. Note that $\tau(v) = y_j$ for some $1 \leq j \leq d$ and that $(\tau')^{-1}(y_0) = v$. If $v \in R$, then τ' satisfies condition (i). Otherwise, $v = x_1$, and τ' satisfies condition (ii) because $y_j \in \tau'(R)$.
- (c) Suppose τ satisfies condition (ii) for some i and neither u nor v equals $\tau^{-1}(y_0)$. Then u and v are contained in the same connected component of $X|_{V(X) \setminus \{x_i\}}$. In particular, we see that

$$\tau'(R \cup \{x_1, \dots, x_{i-1}\}) = \tau(R \cup \{x_1, \dots, x_{i-1}\}),$$

which shows that τ' satisfies condition (ii).

- (d) Suppose τ satisfies condition (ii) for some $2 \leq i \leq d - 1$ and $\tau^{-1}(y_0) \in \{u, v\}$. Without loss of generality, assume $u = \tau^{-1}(y_0) = x_i$. Since the only neighbors of x_i are x_{i-1} and x_{i+1} , we conclude that v is one of these two vertices. Note (as before) that $\tau(v) = y_j$ for some $1 \leq j \leq d$ and that $(\tau')^{-1}(y_0) = v$. If $v = x_{i-1}$, then $(\tau')^{-1}(y_0) = x_{i-1}$ and

$$\tau'(R \cup \{x_1, \dots, x_{i-2}\}) = \tau(R \cup \{x_1, \dots, x_{i-1}\}) \setminus \{y_j\}.$$

The assumption on τ implies that this set has size at least $i - 1$, so τ' satisfies condition (ii). If instead $v = x_{i+1}$, then $(\tau')^{-1}(y_0) = x_{i+1}$ and

$$\tau'(R \cup \{x_1, \dots, x_i\}) = \tau(R \cup \{x_1, \dots, x_{i-1}\}) \cup \{y_j\}.$$

The assumption on τ implies that this set has size at least $i + 1$, so τ' satisfies condition (ii).

- (e) Suppose that $d = 1$, that τ satisfies condition (ii) (necessarily for $i = 1$), and that $\tau^{-1}(y_0) \in \{u, v\}$. Without loss of generality, assume $u = \tau^{-1}(y_0) = x_1$. Note that $v = \tau^{-1}(y_1)$, where $\tau^{-1}(y_1) \in R$ by assumption. Then $(\tau')^{-1}(y_0) = v \in R$, so τ' satisfies condition (i).
- (f) Suppose that $d > 1$, that τ satisfies condition (ii) for $i = 1$, and that $\tau^{-1}(y_0) \in \{u, v\}$. Without loss of generality, assume $u = \tau^{-1}(y_0) = x_1$. Note that we have $\tau(v) = y_j$ for some $1 \leq j \leq d$. Recall that x_1 is adjacent to only x_2 and some vertices in R , so these are the only possible vertices that v can be. If $v = x_2$, then $(\tau')^{-1}(y_0) = x_2$ and

$$\tau'(R \cup \{x_1\}) = \tau(R) \cup \{y_j\}.$$

The assumption on τ implies that this set has size at least 2, so τ' satisfies condition (ii). If $v \in R$, then $(\tau')^{-1}(y_0) = v \in R$, which shows that τ' satisfies condition (i).

- (g) Suppose that $d > 1$, that τ satisfies condition (ii) for $i = d$, and that $\tau^{-1}(y_0) \in \{u, v\}$. Without loss of generality, assume $u = \tau^{-1}(y_0) = x_d$. Note that we have $v = \tau^{-1}(y_j)$ for some $1 \leq j \leq d$. All vertices y_1, \dots, y_d are contained in $\tau(R \cup \{x_1, \dots, x_{d-1}\})$ by assumption, so we see that $v \in R \cup \{x_1, \dots, x_{d-1}\}$. Since v is adjacent to u , we conclude that $v = x_{d-1}$. Then $(\tau')^{-1}(y_0) = x_{d-1}$ and

$$\tau'(R \cup \{x_1, \dots, x_{d-2}\}) = \tau(R \cup \{x_1, \dots, x_{d-1}\}) \setminus \{y_j\}.$$

The assumption on τ implies that this set has size at least $d - 2$, so τ' satisfies condition (ii).

This analysis exhausts the possible cases and completes the proof. □

We now derive some consequences of Theorem 6.1 regarding specific families of graphs. The *lollipop graph* $\text{Lollipop}_{k,m}$ is the graph with vertex set $[k + m]$ and edge set

$$E(\text{Lollipop}_{k,m}) = \{\{i, i + 1\} : 1 \leq i \leq k\} \cup \{\{i, j\} : k + 1 \leq i < j \leq k + m\}.$$

In other words, $\text{Lollipop}_{k,m}$ is obtained by identifying the vertex $k + 1$ in the path graph Path_{k+1} with the vertex $k + 1$ in the complete graph on the vertex set $\{k + 1, \dots, k + m\}$. The following corollary is immediate from Theorem 6.1.

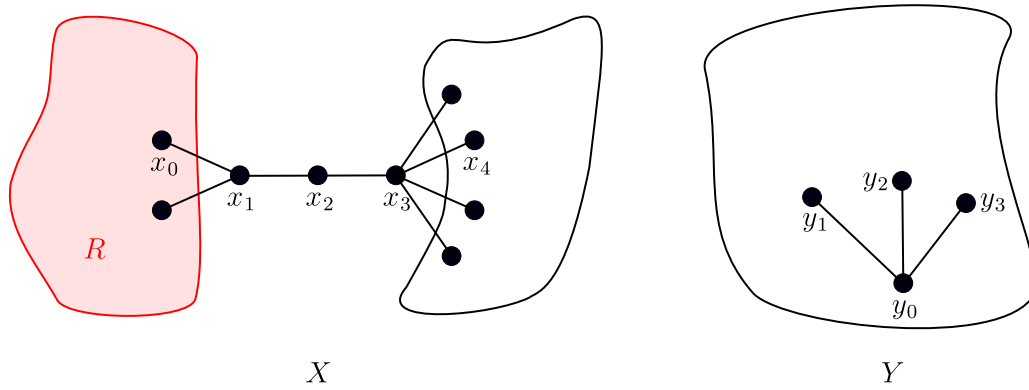


Figure 6.1: A schematic illustration of the proof of Theorem 6.1 (with $d = 3$).

Corollary 6.2. *Let X be a subgraph of $\text{Lollipop}_{n-m,m}$. If Y is a graph on n vertices such that $\text{FS}(X, Y)$ is connected, then the minimum degree of Y is at least $n - m + 1$.*

The graph $\text{Lollipop}_{n-3,3}$ is a prolongation of K_3 . Therefore, Corollary 5.4 guarantees that $\text{FS}(\text{Lollipop}_{n-3,3}, Y)$ is connected whenever Y is a graph on n vertices with minimum degree at least $n - 2$. Consequently, we can appeal to Corollary 6.2 to obtain a complete classification of the graphs Y such that $\text{FS}(\text{Lollipop}_{n-3,3})$ is connected.

Corollary 6.3. *Let Y be a graph on n vertices. The graph $\text{FS}(\text{Lollipop}_{n-3,3}, Y)$ is connected if and only if the minimum degree of Y is at least $n - 2$.*

Remark 6.4. The graph $\text{Lollipop}_{n-5,5}$ is a prolongation of the graph X in Corollary 5.5, so it follows from that corollary that $\text{FS}(\text{Lollipop}_{n-5,5}, Y)$ is connected whenever Y is a graph on n vertices with minimum degree at least $n - 3$. On the other hand, Corollary 6.2 tells us that $\text{FS}(\text{Lollipop}_{n-5,5}, Y)$ is disconnected whenever Y is a graph on n vertices with minimum degree at most $n - 5$. We leave open the characterization of graphs Y such that $\text{FS}(\text{Lollipop}_{n-5,5}, Y)$ is connected; in order to finish the characterization, it suffices to consider only the graphs Y with minimum degree $n - 4$.

We end this section with a discussion of one additional nice family of graphs. For $n \geq 3$, let D_n be the graph with vertex set $[n]$ and edge set $\{\{i, i+1\} : 1 \leq i \leq n-2\} \cup \{\{n-2, n\}\}$. The notation D_n is chosen because this graph is the Dynkin diagram of type D_n . We will show that for $n \geq 4$, the characterization of the graphs Y such that $\text{FS}(D_n, Y)$ is connected is precisely the same as the characterization of the graphs Y such that $\text{FS}(\text{Lollipop}_{n-3,3}, Y)$ is connected. Hence, the edge $\{n-1, n\}$ in $\text{Lollipop}_{n-3,3}$ has no influence on whether or not $\text{FS}(\text{Lollipop}_{n-3,3}, Y)$ is connected.

Theorem 6.5. *Let Y be a graph on $n \geq 4$ vertices. The graph $\text{FS}(D_n, Y)$ is connected if and only if the minimum degree of Y is at least $n - 2$.*

Before proving this result, we establish a simple lemma. Recall that $\alpha_{\bar{Y}}(\lambda)$ denotes the acyclic orientation of \bar{Y} associated to the permutation λ , where an edge $\{y, z\}$ is oriented from y to z if and only if $\lambda^{-1}(y) < \lambda^{-1}(z)$.

Lemma 6.6. *Let Y be a graph with vertex set $[n]$. If λ and λ' are two permutations in \mathfrak{S}_n such that $\lambda(n) = \lambda'(n)$ and $\alpha_{\overline{Y}}(\lambda) = \alpha_{\overline{Y}}(\lambda')$, then λ and λ' are in the same connected component of $\text{FS}(\text{D}_n, Y)$.*

Proof. Let $Y' = Y|_{[n] \setminus \{\lambda(n)\}}$. The restrictions $\lambda|_{[n-1]}$ and $\lambda'|_{[n-1]}$ are vertices of $\text{FS}(\text{Path}_{n-1}, Y')$ such that $\alpha_{\overline{Y'}}(\lambda|_{[n-1]}) = \alpha_{\overline{Y'}}(\lambda'|_{[n-1]})$, so it follows immediately from Theorem 3.1 that there is a sequence of (Path_{n-1}, Y') -friendly swaps transforming $\lambda|_{[n-1]}$ into $\lambda'|_{[n-1]}$. Because Path_{n-1} is a subgraph of D_n and Y' is a subgraph of Y , this same sequence of friendly swaps can be viewed as a sequence of (D_n, Y) -friendly swaps that transforms λ into λ' . \square

Proof of Theorem 6.5. We may assume that $V(Y) = [n]$. Observe that the graph D_n is a subgraph of $\text{Lollipop}_{n-3,3}$. If the minimum degree of Y is at most $n - 3$, then it follows from Proposition 2.1 and Corollary 6.3 that $\text{FS}(\text{D}_n, Y)$ is disconnected.

Now assume Y has minimum degree at least $n - 2$. This means that \overline{Y} has maximum degree at most 1. In other words, \overline{Y} is a disjoint union of copies of K_1 and K_2 . Choose a permutation $\sigma \in \mathfrak{S}_n$, and let B be the connected component of $\text{FS}(\text{D}_n, Y)$ containing σ . We are going to show that the identity permutation in \mathfrak{S}_n is also in B . Since σ was arbitrary, this will imply that $\text{FS}(\text{D}_n, Y)$ is connected. Consider the acyclic orientation $\alpha_{\overline{Y}}(\sigma)$ of \overline{Y} associated to σ . Suppose $\{u, v\}$ is an edge of \overline{Y} with $\sigma^{-1}(u) < \sigma^{-1}(v)$. This means that in $\alpha_{\overline{Y}}(\sigma)$, the edge $\{u, v\}$ is oriented from u to v . Since $\{u, v\}$ is an isolated edge in \overline{Y} , the vertex v must be a sink of $\alpha_{\overline{Y}}(\sigma)$. We claim that there is a permutation $\hat{\sigma}$ in B such that $\alpha_{\overline{Y}}(\hat{\sigma})$ is the same as $\alpha_{\overline{Y}}(\sigma)$ except with the orientation of $\{u, v\}$ reversed. First, suppose $\sigma(n) = v$. Let r be such that $\sigma(n-r) = u$. We can apply (D_n, Y) -friendly swaps across $\{n-r, n-r+1\}, \{n-r+1, n-r+2\}, \dots, \{n-2, n-1\}$ (in this order) and then another (D_n, Y) -friendly swap across $\{n-2, n\}$ so as to obtain the permutation

$$\hat{\sigma} = \sigma \circ (n-r \ n-r+1) \circ (n-r+1 \ n-r+2) \circ \dots \circ (n-2 \ n-1) \circ (n-2 \ n).$$

It is straightforward to check that $\hat{\sigma}$ is in B and that $\alpha_{\overline{Y}}(\hat{\sigma})$ is obtained from $\alpha_{\overline{Y}}(\sigma)$ by reversing the orientation of $\{u, v\}$. Next, assume that $\sigma(n) \neq v$. Since u is a source of $\alpha_{\overline{Y}}(\sigma)$, we also have $\sigma(n) \neq u$. There exists a permutation $\sigma' \in \mathfrak{S}_n$ such that $\alpha_{\overline{Y}}(\sigma') = \alpha_{\overline{Y}}(\sigma)$, $\sigma'(n) = \sigma(n)$, $\sigma'(n-1) = v$, and $\sigma'(n-2) = u$. Applying Lemma 6.6 with $\lambda = \sigma$ and $\lambda' = \sigma'$, we find that σ' is in B . If we apply a (D_n, Y) -friendly swap across $\{n-2, n\}$ to σ' , we obtain a permutation $\hat{\sigma}$ in B such that the acyclic orientations $\alpha_{\overline{Y}}(\sigma)$ and $\alpha_{\overline{Y}}(\hat{\sigma})$ agree except in the orientation of $\{u, v\}$, as desired.

We have shown that we can use a sequence of (D_n, Y) -friendly swaps to reverse the orientation of a single edge in the acyclic orientation of \overline{Y} associated to a permutation. By repeating this argument, we can eventually construct a permutation τ in B such that $\alpha_{\overline{Y}}(\tau)$ is any prescribed acyclic orientation of \overline{Y} . In particular, we can choose τ in B so that each edge $\{a, b\}$ in \overline{Y} with $a < b$ is oriented from a to b in $\alpha_{\overline{Y}}(\tau)$. In other words, $\alpha_{\overline{Y}}(\tau) = \beta$, where β is the acyclic orientation of \overline{Y} associated to the identity permutation in \mathfrak{S}_n .

We are going to show that there is a permutation τ' in B with $\alpha_{\overline{Y}}(\tau') = \beta$ and $\tau'(n) = n$. If we can do this, then it will follow from Lemma 6.6 (with $\lambda = \tau'$ and λ' the identity permutation) that the identity permutation is also in B , as claimed. If $\tau(n) = n$, then we can simply set $\tau' = \tau$. Thus, let us assume that $\tau(n) \neq n$. Using the fact that $n \geq 4$, it is not difficult to check

that there exists a linear extension $\widehat{\tau}$ of β such that $\widehat{\tau}(n) = \tau(n)$ and $\widehat{\tau}(n-2) = n$. Invoking Lemma 6.6 with $\lambda = \tau$ and $\lambda' = \widehat{\tau}$, we see that $\widehat{\tau}$ is in B . Note that n is a sink of β by the definition of β . Since $\widehat{\tau}$ is a linear extension of β , the vertex $\widehat{\tau}(n)$ must also be a sink of β . The vertices $\widehat{\tau}(n)$ and n are distinct sinks of β , so they are not adjacent in \overline{Y} . Therefore, we can perform a (D_n, Y) -friendly swap across $\{n-2, n\}$ in order to transform $\widehat{\tau}$ into a different linear extension τ' of β with $\tau'(n) = n$. This permutation τ' is in B , as desired. \square

7. Concluding Remarks and Open Problems

We conclude with several open problems and suggestions for future inquiries along the lines of the present paper.

7.1. Other graphs

In Sections 3 and 4, we gained a full understanding of the connected components of $\text{FS}(\text{Path}_n, Y)$ and $\text{FS}(\text{Cycle}_n, Y)$. It could be interesting consider complements of paths and cycles by investigating the connected components of $\text{FS}(\overline{\text{Path}}_n, Y)$ and $\text{FS}(\overline{\text{Cycle}}_n, Y)$. Another natural direction would be the exploration of graphs of the form $\text{FS}(K_{k, n-k}, Y)$, where $K_{k, n-k}$ is the complete bipartite graph with partite sets of sizes k and $n-k$. Note that $K_{1, n-1}$ is isomorphic to Star_n , which was studied thoroughly by Wilson [23]. It might be interesting to consider just the specific complete bipartite graphs $K_{2, n-2}$, or possibly $K_{r, r}$. It would also be interesting to obtain more general results about the graphs $\text{FS}(X, Y)$ when X is a tree (or even a specific type of tree). Let us also recall Remark 6.4, which asks for a characterization of the graphs Y such that $\text{FS}(\text{Lollipop}_{n-5, 5}, Y)$ is connected. Of course, several other nice families of graphs that we have not mentioned could also give rise to interesting results.

7.2. Making $\text{FS}(X, Y)$ connected for all reasonable graphs Y

One might naturally ask for the sparsest graph Y on n vertices such that $\text{FS}(X, Y)$ is connected whenever X is a connected graph on n vertices. For $\text{FS}(\text{Path}_n, Y)$ to be connected, Y must be a complete graph. A less trivial variant of this problem asks for the sparsest graph Y on n vertices such that $\text{FS}(X, Y)$ is connected whenever X is a biconnected graph on n vertices. In this case, $\text{FS}(\text{Cycle}_n, Y)$ must be connected, so it follows from Corollary 4.14 that \overline{Y} must be a forest consisting of trees of coprime sizes. This implies that Y must have at least $\binom{n}{2} - (n-2)$ edges.

This bound of $\binom{n}{2} - (n-2)$ is tight. Indeed, suppose Y is a graph such that \overline{Y} is the disjoint union of a tree on $n-1$ vertices and an isolated vertex v^* . Then Y contains a subgraph isomorphic to Star_n (with v^* as the center of the star). Suppose X is biconnected. We want to show that $\text{FS}(X, Y)$ is connected. If X is isomorphic to Cycle_n , then the connectedness of $\text{FS}(X, Y)$ follows from Corollary 4.14. If $n = 7$ and X is isomorphic to the graph θ_0 from Theorem 2.5, then one can check by computer that $\text{FS}(X, Y)$ is connected. Now assume X is not isomorphic to θ_0 or a cycle graph. If X is not bipartite, then $\text{FS}(X, Y)$ is connected because $\text{FS}(X, \text{Star}_n)$ is connected by Wilson's theorem ($\text{FS}(X, \text{Star}_n)$ and $\text{FS}(\text{Star}_n, X)$ are

isomorphic). If X is bipartite, then the connectedness of $\text{FS}(X, Y)$ follows from Remark 2.8 because Star_n is isomorphic to a proper subgraph of Y .

We have shown that if $\text{FS}(X, Y)$ is connected for all biconnected graphs X , then \bar{Y} must be a forest consisting of trees of coprime sizes. We have seen that some of these choices for Y do indeed make $\text{FS}(X, Y)$ connected for all biconnected X . The following conjecture states that *all* such choices of Y satisfy this property.

Conjecture 7.1. Let Y be a graph on n vertices such that \bar{Y} is a forest consisting of trees $\mathcal{T}_1, \dots, \mathcal{T}_r$ such that $\gcd(|V(\mathcal{T}_1)|, \dots, |V(\mathcal{T}_r)|) = 1$. If X is a biconnected graph on n vertices, then $\text{FS}(X, Y)$ is connected.

7.3. Diameter and girth

In a different direction, one might ask about the diameter and girth of the graph $\text{FS}(X, Y)$ or its connected components. Recall that if X and Y are graphs on n vertices, then $\text{FS}(X, Y)$ has $n!$ vertices. Must the diameter of a connected component of $\text{FS}(X, Y)$ be polynomially bounded in n ? What changes if we restrict our attention to X and Y such that $\text{FS}(X, Y)$ is connected? For the girth of $\text{FS}(X, Y)$, we can say somewhat more. Since $\text{FS}(X, Y)$ is bipartite (by Proposition 2.2), its girth is either even (and greater than or equal to 4) or infinite. It is easy to see that if X and Y each have 2 pairwise disjoint edges, then the girth of $\text{FS}(X, Y)$ is exactly 4. The behavior in the remaining case (which can be reduced to the setting where X is a star and Y is connected) remains unknown.

It would also be interesting to understand the diameters of connected components of some specific graphs that we have studied already, such as $\text{FS}(\text{Path}_n, Y)$ and $\text{FS}(\text{Cycle}_n, Y)$. For example, Proposition 1.3 states that for every n -element poset P , the group of bijections generated by the toggle operators t_1, \dots, t_{n-1} acts transitively on the set $\mathcal{L}(P)$ of linear extensions of P . It is natural to ask for the maximum distance between two linear extensions of P , where the distance is measured by the number of toggle operators needed to change the first linear extension into the second. To state this problem in the language we have been using throughout this article, let Z be the Hasse diagram of P (viewed as a graph), and let $Y = \bar{Z}$. Consider the acyclic orientation α of $\bar{Y} = Z$ given by orienting each edge $\{x, y\}$ from x to y whenever y covers x in P . Then we are asking for the diameter of the connected component H_α of $\text{FS}(\text{Path}_n, Y)$ as defined in Theorem 3.1.

7.4. New equivalence relations for acyclic orientations

Equivalences of acyclic orientations under various flips played a major role in our analysis in Section 4. Fix nonnegative integers a and b and a graph G , and consider the set $\text{Acyc}(G)$ of acyclic orientations of G . Suppose that in an acyclic orientation $\alpha \in \text{Acyc}(G)$, there are pairwise non-adjacent vertices u_1, \dots, u_a and v_1, \dots, v_b such that every u_i is a source and every v_i is a sink. Then we can obtain a new acyclic orientation by reversing the directions of all edges incident to any of these vertices; equivalently, we simultaneously flip every v_i into a sink and every u_i into a source. Call such an operation an (a, b) -flip. We say that acyclic orientations α, α' are $\{a, b\}$ -flip equivalent if α' can be obtained from α by a sequence of (a, b) -flips and

(b, a) -flips. (The reader can easily verify that this is in fact an equivalence relation.) In this language, the toric acyclic orientations discussed thoroughly in Section 4 (which have received significant attention in previous articles such as [10]) are simply $\{0, 1\}$ -flip equivalence classes. Furthermore, the double-flip equivalence classes introduced in Section 4, which are used to parameterize the connected components of graphs of the form $\text{FS}(\text{Cycle}_n, Y)$, are simply $\{1, 1\}$ -flip equivalence class. We also saw that there is a nontrivial connection between $\{0, 1\}$ -flip equivalence classes and $\{1, 1\}$ -flip equivalence classes: if G has connected components of sizes n_1, \dots, n_r and $\nu = \gcd(n_1, \dots, n_r)$, then every $\{0, 1\}$ -flip equivalence class is the union of exactly ν $\{1, 1\}$ -flip equivalence classes. Furthermore, we saw that for each fixed $\{0, 1\}$ -flip equivalence class $[\alpha]_{\sim}$ of $\text{Acyc}(G)$, the ν $\{1, 1\}$ -flip equivalence classes contained in $[\alpha]_{\sim}$ all have the same number of linear extensions.

What can be said about $\{a, b\}$ -flip equivalence classes for other choices of $\{a, b\}$? Are there other interesting connections between $\{a, b\}$ -flip equivalence classes and $\{a', b'\}$ -flip equivalence classes for different choices of a, a', b, b' ?

7.5. Probabilistic and extremal questions

It is natural to ask about the connectedness of the friends-and-strangers graph $\text{FS}(X, Y)$ that results from choosing X and Y to be independent Erdős-Rényi random graphs in $\mathcal{G}(n, p)$. In a separate article with Noga Alon [1], we prove that the threshold probability p where $\text{FS}(X, Y)$ changes from being disconnected with high probability to being connected with high probability is $p = n^{-1/2+o(1)}$.

In the same article, we obtain estimates for the minimum d such that whenever X and Y are n -vertex graphs with minimum degrees at least d , the friends-and-strangers graph $\text{FS}(X, Y)$ is connected.

7.6. Right versus left multiplication

One of the referees suggested looking at analogues of friends-and-strangers graphs for other Coxeter groups. The construction discussed in this section provides one such analogue that works for any group.

Suppose Γ is a group and $S \subseteq \Gamma$ is closed under taking inverses, meaning $\tau^{-1} \in S$ for every $\tau \in S$. The *right Cayley graph*³ $\text{Cay}_R(\Gamma, S)$ is the graph with vertex set Γ and edge set

$$E(\text{Cay}_R(\Gamma, S)) = \{\{\sigma, \sigma\tau\} : \sigma \in \Gamma, \tau \in S\}.$$

In other words, edges in $\text{Cay}_R(\Gamma, S)$ correspond to right multiplication by elements of S . Similarly, the *left Cayley graph* $\text{Cay}_L(\Gamma, S)$ is the graph with vertex set \mathfrak{S}_n and edge set

$$E(\text{Cay}_L(\Gamma, S)) = \{\{\sigma, \tau\sigma\} : \sigma \in \Gamma, \tau \in S\}.$$

For $i, j \in [n]$, let us identify the pair $\{i, j\}$ with the transposition (ij) in \mathfrak{S}_n that swaps i and j . If X and Y are graphs with $V(X) = V(Y) = [n]$, then the edge set of the friends-and-strangers graph of X and Y is

$$E(\text{FS}(X, Y)) = E(\text{Cay}_R(\mathfrak{S}_n, E(X))) \cap E(\text{Cay}_L(\mathfrak{S}_n, E(Y))).$$

³This deviates slightly from standard terminology since we do not require S to generate Γ .

In particular, we have $\text{FS}(X, K_n) = \text{Cay}_R(\mathfrak{S}_n, E(X))$ and $\text{FS}(K_n, Y) = \text{Cay}_L(\mathfrak{S}_n, E(Y))$. This suggests a natural generalization of friends-and-strangers graphs. Namely, if S_R and S_L are subsets of a group Γ that are each closed under taking inverses, then we can consider the graph with vertex set Γ and edge set

$$E(\text{Cay}_R(\Gamma, S_R)) \cap E(\text{Cay}_L(\Gamma, S_L)).$$

The preceding definition seems far too broad for one to say anything meaningful in full generality; however, it is possible that there could be interesting directions to explore for specific groups (for instance, following the referee's suggestion, hyperoctahedral groups). Even when $\Gamma = \mathfrak{S}_n$, the above definition still vastly generalizes friends-and-strangers graphs because we do not require the elements of S_R and S_L to be transpositions. It could be fruitful to investigate whether these graphs have interesting structural properties when $\Gamma = \mathfrak{S}_n$ and the sets S_R and S_L are chosen very specifically.

Acknowledgements

We thank Christian Gaetz for pointing us to the article [22] and Alexander Postnikov for showing us that the connected components of $\text{FS}(\text{Path}_n, Y)$ correspond to the acyclic orientations of \bar{Y} . We thank Noga Alon and Matthew Macauley for engaging in very helpful conversations. We also thank the anonymous referees for their thorough reading and detailed comments. We are especially grateful to the referee who outlined their own very detailed alternative view on friends-and-strangers graphs.

References

- [1] N. Alon, C. Defant, and N. Kravitz, Typical and extremal aspects of friends-and-strangers graphs. *ArXiv preprint*, arXiv:2009.07840v1 (2020).
- [2] C. Barrett and C. Reidys, Elements of a theory of computer simulation I: sequential CA over random graphs. *Appl. Math. Comp.*, **98** (1999), 241–259.
- [3] A. Barvinok, Matrices with prescribed row and column sums. *Linear Algebra Appl.*, **436** (2012), 820–844.
- [4] A. Biniiaz, K. Jain, A. Lubiw, Z. Masárová, T. Miltzow, D. Mondal, A. Naredla, J. Tkadlec, and A. Turcotte, Token swapping on trees. *ArXiv preprint*, arXiv:1903.06981v2 (2019).
- [5] A. Björner and M. Wachs, Permutation statistics and linear extensions of posets. *J. Combin. Theory Ser. A*, **58** (1991), 85–114.
- [6] B. Bollobás, L. Pebody, and O. Riordan, Contraction-deletion invariants for graphs. *J. Combin. Theory Ser. B*, **80** (2000), 320–345.
- [7] P. Cartier and D. Foata, Problèmes combinatoires de commutation et réarrangements. *Lect. Notes Math.* **85**, Springer–Verlag, Berlin–New York 1969.
- [8] B. Chen, Orientations, lattice polytopes, and group arrangements I, chromatic and tension polynomials of graphs. *Ann. Comb.*, **13** (2010), 425–452.

- [9] J. Conway, N. Elkies, and J. Martin, The Mathieu group M_{12} and its pseudogroup extension M_{13} . *Experiment. Math.*, **15.2** (2006), 223–236.
- [10] M. Develin, M. Macauley, and V. Reiner, Toric partial orders. *Trans. Amer. Math. Soc.*, **368** (2016), 2263–2287.
- [11] P. Diaconis and A. Gangolli, Rectangular arrays with fixed margins. *Discrete Probability and Algorithms* (New York: Springer, 1995), 15–41.
- [12] H. Eriksson and K. Eriksson, Conjugacy of Coxeter elements. *Electron. J. Combin.*, **16** (2009).
- [13] N. Gill, N. Gillespie, A. Nixon, and J. Semeraro, Generating groups using hypergraphs. *Q. J. Math.*, **67** (2016), 29–52.
- [14] C. Greene and T. Zaslavsky, On the interpretation of Whitney numbers through arrangements of hyperplanes, zonotopes, non-Radon partitions, and orientations of graphs. *Trans. Amer. Math. Soc.*, **280** (1982), 97–126.
- [15] T. Lam and L. Williams, Total positivity for cominiscule Grassmannians. *New York J. Math.*, **14** (2008), 53–99.
- [16] M. Macauley and H. S. Mortveit, Posets from admissible Coxeter sequences. *Electron. J. Combin.*, **18** (2011).
- [17] P.-L. Méliot, *Representation theory of symmetric groups*. CRC Press, 2017.
- [18] K. M. Mosesjan, Strongly basable graphs. *Akad. Nauk Armjan. SSR Dokl.*, **54** (1972), 134–138. (Russian, with Armenian summary.)
- [19] I. Pak, Hook length formula and geometric combinatorics. *Sém. Lothar. Combin.*, **46** (2001/2).
- [20] O. Pretzel, On reorienting graphs by pushing down maximal vertices. *Order*, **3** (1986), 135–153.
- [21] D. E. Speyer, Powers of Coxeter elements in infinite groups are reduced. *Proc. Amer. Math. Soc.*, **137** (2009), 1295–1302.
- [22] R. P. Stanley, An equivalence relation on the symmetric group and multiplicity-free flag h -vectors. *J. Comb.*, **3** (2012), 277–298.
- [23] R. M. Wilson, Graph puzzles, homotopy, and the alternating group. *J. Combin. Theory, Ser. B*, **16** (1974), 86–96.
- [24] C. Yang, Sliding puzzles and rotating puzzles on graphs. *Discrete Math.*, **311** (2011), 1290–1294.