Séminaire Lotharingien de Combinatoire **85B** (2021)
Article #21, 12 pp.

Proceedings of the 33rd Conference on Formal Power Series and Algebraic Combinatorics (Ramat Gan)

Friends and strangers walking on graphs

Noga Alon*[1,2], Colin Defant†[1], and Noah Kravitz‡[1]

[1]*Department of Mathematics, Princeton University, Princeton, NJ 08544, USA*
[2]*Schools of Mathematics and Computer Science, Tel Aviv University, Tel Aviv 69978, Israel*

Abstract. We introduce the friends-and-strangers graph $\mathsf{FS}(X,Y)$ associated with graphs X and Y whose vertex sets $V(X)$ and $V(Y)$ have the same cardinality. This is the graph whose vertex set consists of all bijections $\sigma : V(X) \to V(Y)$, 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. The most fundamental questions that one can ask about these friends-and-strangers graphs concern their connected components and, in particular, when there is only a single connected component.

When X is a path graph, we show that the connected components of $\mathsf{FS}(X,Y)$ 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 $\mathsf{FS}(X,Y)$ in terms of toric acyclic orientations of the complement of Y. In a more probabilistic vein, we address the case of "typical" X and Y by proving that if X and Y are independent Erdős-Rényi random graphs with n vertices and edge probability p, then the threshold probability guaranteeing the connectedness of $\mathsf{FS}(X,Y)$ with high probability is $p = n^{-1/2+o(1)}$. We also study the case of "extremal" X and Y by proving that the smallest minimum degree of the n-vertex graphs X and Y that guarantees the connectedness of $\mathsf{FS}(X,Y)$ is between $3n/5 + O(1)$ and $9n/14 + O(1)$. Furthermore, we obtain bipartite analogues of the latter two results.

Keywords: friends-and-strangers graph, acyclic orientation, toric acyclic orientation

1 Introduction

Given simple graphs X and Y on n vertices, we define the *friends-and-strangers graph* of X and Y, denoted $\mathsf{FS}(X,Y)$, as follows. The vertex set of $\mathsf{FS}(X,Y)$ is the set of all bijections

*nalon@math.princeton.edu. The author is supported by NSF grant DMS-1855464, BSF grant 2018267, and the Simons Foundation.

†cdefant@princeton.edu. The author is supported by a Fannie and John Hertz Foundation Fellowship and NSF Graduate Research Fellowship grant DGE-1656466.

‡nkravitz@princeton.edu.

$\sigma : V(X) \to V(Y)$ from the vertex set of X to the vertex set of Y; two bijections σ and σ' are connected by an edge if and only if X contains an edge $\{a, b\}$ 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 $c \in V(X) \setminus \{a, b\}$. In other words, we connect σ and σ' if they differ only at a pair of adjacent vertices such that the images of these vertices under σ are adjacent in Y. In this case, the operation that transforms σ into σ' is called an (X, Y)-*friendly swap*.

The friends-and-strangers graph $\mathsf{FS}(X, Y)$ has the following non-technical interpretation. Identify n different people with the vertices of Y. Say that two such people are friends with each other if they are adjacent in Y, and say that they are strangers otherwise. Now, suppose that these people are standing on the vertices of X so that each vertex has exactly one person standing on it. At each point in time, two friends standing on adjacent vertices of X may swap places, but two strangers may not. It is natural to ask how various configurations can be reached from other configurations when we allow multiple such swaps, and this is precisely the information that is encoded in $\mathsf{FS}(X, Y)$. In particular, the connected components of $\mathsf{FS}(X, Y)$ correspond to the equivalence classes of mutually-reachable configurations.

It is sometimes convenient to assume that $V(X)$ and $V(Y)$ are both the set $[n] := \{1, \ldots, n\}$. In this case, the vertices of $\mathsf{FS}(X, Y)$ are the elements of the symmetric group $\mathfrak{S}_n$, which consists of all permutations of the numbers $1, \ldots, n$. For $i, j \in [n]$, let $(i\,j)$ 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 $\sigma \circ (i\,j)$. If we write the permutation σ in one-line notation as $\sigma = \sigma(1) \cdots \sigma(n)$, 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 $X = Y =$ 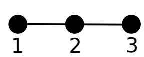

, then $\mathsf{FS}(X, Y) =$ 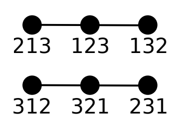

.

This framework is quite general, and several special cases have received attention in the past in other contexts. For instance, Stanley [7] studied the connected components of $\mathsf{FS}(\mathsf{Path}_n, \mathsf{Path}_n)$; the graph $\mathsf{FS}(K_n, Y)$ is the Cayley graph of $\mathfrak{S}_n$ generated by the transpositions corresponding to edges of Y; the famous 15-puzzle can be interpreted in terms of $\mathsf{FS}(X, Y)$ where X is the star graph on 16 vertices and Y is the 4×4 grid graph; and Wilson [8], generalizing the 15-puzzle, characterized the graphs Y such that $\mathsf{FS}(X, Y)$ is connected when X is a star graph.

This extended abstract is based on the articles [5] and [1]: the first establishes several general properties of friends-and-strangers graphs and characterizes the connected components of $\mathsf{FS}(\mathsf{Path}_n, Y)$ and $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ for arbitrary graphs Y; the second studies the behavior of $\mathsf{FS}(X, Y)$ for random graphs X and Y and addresses a few extremal questions related to the connectivity of $\mathsf{FS}(X, Y)$.

We will explain how the connected components of $\mathsf{FS}(\mathsf{Path}_n, Y)$ are parameterized by the acyclic orientations of the complement of Y; this result is essentially equivalent to the well-known fact that the Coxeter elements of a Coxeter system correspond to the acyclic orientations of the Coxeter graph. We will also explain how the connected components of $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ are closely related to toric acyclic orientations (also called toric partial orders), which have appeared in many contexts and were formalized in [6]. The connected components of $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ can be understood via a new equivalence relation on acyclic orientations of the complement of Y that we call double-flip equivalence; this new notion could be of independent interest. It turns out that our analysis of the graphs $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ 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 $\mathsf{FS}(\mathsf{Cycle}_n, Y)$, where ν is the greatest common divisor of the sizes of the connected components of the complement of Y. One corollary is that $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ is connected if and only if the complement of Y is a forest whose trees have coprime sizes. We will also briefly mention some results concerning when $\mathsf{FS}(X, Y)$ is connected for some other specific choices of X.

It is natural to ask about the connected components of $\mathsf{FS}(X, Y)$ when X and Y are Erdős-Rényi random edge-subgraphs of the complete graph K_n with edge probability p. We find that $p = n^{-1/2+o(1)}$ is the threshold at which $\mathsf{FS}(X, Y)$ changes from disconnected with high probability to connected with high probability. We also obtain an analogous (but less tight) result for random bipartite graphs X and Y; in this case, a simple parity obstruction prevents $\mathsf{FS}(X, Y)$ from being connected, so we examine when $\mathsf{FS}(X, Y)$ has exactly 2 connected components.

Next, from a more extremal point of view, we find that the smallest minimum degree of the n-vertex graphs X and Y that guarantees the connectedness of $\mathsf{FS}(X, Y)$ is between $3n/5 + O(1)$ and $9n/14 + O(1)$. In the analogous bipartite setting where X and Y are edge-subgraphs of the complete bipartite graph $K_{r,r}$, we find that the cutoff minimum degree for having exactly 2 connected components is either $\lceil (3r+1)/4 \rceil$ or $\lceil (3r+2)/4 \rceil$.

Because the study of friends-and-strangers graphs is quite young, there remain many promising open questions; we raise a few at the end of this extended abstract.

1.1 Notation and terminology

In what follows, let $V(G)$ and $E(G)$ denote the vertex set and edge set (respectively) of a graph G. Some specific graphs with vertex set $[n]$ that will play a large role for us are: the *path graph* Path_n, which has edge set $E(\mathsf{Path}_n) = \{\{i, i+1\} : i \in [n-1]\}$; and the *cycle graph* Cycle_n, which has edge set $E(\mathsf{Cycle}_n) = \{\{i, i+1\} : i \in [n-1]\} \cup \{\{n, 1\}\}$.

The *complement* of a graph 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 $\{a, b\} \notin$

$E(G)$. Also, 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 $\bigoplus_{i \in I} G_i$.

We obtain an *orientation* of a graph G by choosing a direction for each of its edges. An orientation is *acyclic* if it does not contain a directed cycle. Let $\mathsf{Acyc}(G)$ denote the set of acyclic orientations of G. If $V(G) = [n]$, then we obtain from each $\alpha \in \mathsf{Acyc}(G)$ 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 $([n], \leq_\alpha)$, where a *linear extension* of the poset $([n], \leq_P)$ is a permutation $\sigma \in \mathfrak{S}_n$ such that $\sigma^{-1}(a) \leq \sigma^{-1}(b)$ whenever $a \leq_P b$. For each permutation $\sigma \in \mathfrak{S}_n$, there is a unique acyclic orientation $\alpha_G(\sigma) \in \mathsf{Acyc}(G)$ 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 $\sigma^{-1}(i) < \sigma^{-1}(j)$.

2 Paths

The main purpose of this section is to discuss the following theorem, which states that the connected components of $\mathsf{FS}(\mathsf{Path}_n, Y)$ correspond to the acyclic orientations of $\overline{Y}$ and that the vertices within each connected component are the linear extensions of the corresponding acyclic orientation.

Theorem 2.1 ([5])**.** *Let Y be a graph with vertex set $[n]$. For each $\alpha \in \mathsf{Acyc}(\overline{Y})$, choose a linear extension $\sigma_\alpha \in \mathcal{L}(\alpha)$, and let H_α be the connected component of $\mathsf{FS}(\mathsf{Path}_n, Y)$ containing σ_α. The connected component H_α depends only on α (not on the specific choice of σ_α), and its vertex set is $\mathcal{L}(\alpha)$. Moreover,*

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

Recall that a *Coxeter system* is a pair (W, S), where W is a group with generating set $S = \{s_1, \ldots, s_n\}$ and presentation $W = \langle S : (s_i s_j)^{m_{i,j}} = 1 \rangle$. Here, the exponents $m_{i,j}$ are elements of $\{1, 2, 3 \ldots\} \cup \{\infty\}$ such that $m_{ii} = 1$ for all $i \in [n]$ and $m_{i,j} \geq 2$ 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* 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* of (W, S) is an element of W of the form $s_{\sigma(1)} \cdots s_{\sigma(n)}$, where $\sigma \in \mathfrak{S}_n$.

Now let Y be a graph with vertex set $[n]$. There exists a Coxeter system (W, S) whose Coxeter graph is $\overline{Y}$, where we identify the vertex $i \in [n] = V(\overline{Y})$ with the element $s_i \in S$. With this identification, every permutation $\sigma \in \mathfrak{S}_n$ gives rise to a word $s_{\sigma(1)} \cdots s_{\sigma(n)}$, which in turn 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 $s_{\sigma(1)} \cdots s_{\sigma(n)}$ 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 $s_{\sigma(1)} \cdots s_{\sigma(n)}$ and $s_{\sigma'(1)} \cdots s_{\sigma'(n)}$ are equal if and only if σ and σ' are in the same connected component of $\mathsf{FS}(\mathsf{Path}_n, Y)$. It follows that Theorem 2.1 is equivalent to the following standard theorem about Coxeter elements.

Theorem 2.2 ([4, 6])**.** *Let (W, S) be a Coxeter system with Coxeter graph G, and write $S = \{s_1, \ldots, s_n\}$. Identify each vertex s_i of G with the element i of $[n]$. For each acyclic orientation $\alpha \in \mathsf{Acyc}(G)$, choose a linear extension σ_α of $([n], \leq_\alpha)$. The Coxeter element $s_{\sigma_\alpha(1)} \cdots s_{\sigma_\alpha(n)}$ depends only on α, not on the specific linear extension σ_α. Furthermore, the map $\alpha \mapsto s_{\sigma_\alpha(1)} \cdots s_{\sigma_\alpha(n)}$ is a bijection from $\mathsf{Acyc}(G)$ to the set of Coxeter elements of (W, S).*

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. Furthermore, a graph with at least 1 edge has at least 2 acyclic orientations. Hence, we have the following corollary.

Corollary 2.3 ([5])**.** *Let Y be a graph with vertex set $[n]$. The number of connected components of $\mathsf{FS}(\mathsf{Path}_n, Y)$ is $T_{\overline{Y}}(2, 0)$. In particular, $\mathsf{FS}(\mathsf{Path}_n, Y)$ is connected if and only if $Y = K_n$.*

3 Cycles

In this section, we describe the structure of $\mathsf{FS}(\mathsf{Cycle}_n, Y)$. The interpretation of people standing on the vertices of a graph is especially natural in this case: the vertices of Y represent people, the edges of Y represent friendships, the people stand around in a circle, and two friends are allowed to swap places if they are standing next to each other in the circle. As a specific corollary of our results in this section (Corollary 3.4), we will see that we can get from any configuration of people standing around the circle to any other configuration if and only if $\overline{Y}$ is a forest whose trees are of relatively prime sizes. 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. However, the proofs (which we omit in this extended abstract) are much more involved.

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 in α; 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 of the edges incident to v. We call this operation a *flip*. Two acyclic orientations $\alpha, \alpha' \in \mathsf{Acyc}(G)$ are *torically equivalent*, denoted $\alpha \sim \alpha'$, if α' can be obtained from α via a sequence of flips (and this is easily seen to be an equivalence relation). The equivalence classes in $\mathsf{Acyc}(G)/\sim$ are called *toric acyclic orientations* (or sometimes *toric partial orders*). We denote the toric acyclic orientation containing the acyclic orientation α by $[\alpha]_\sim$.

Toric acyclic orientations have been studied in many different forms (see [6] and other references in [5]); the article [6] formalizes a systematic framework for their investigation. One of the reasons for the use of the word "toric" stems from a connection with hyperplane arrangements. Indeed, there is a natural one-to-one correspondence between the connected components of $(\mathbb{R}^n/\mathbb{Z}^n) \setminus \pi(\mathcal{A}(G))$ and the toric acyclic orientations of G, where $\mathcal{A}(G)$ is the graphical hyperplane arrangement of the graph G and $\pi : \mathbb{R}^n \to \mathbb{R}^n/\mathbb{Z}^n$ is the natural projection.

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 $\varphi : [n] \to [n]$ be the cyclic permutation given by $\varphi(i) = i + 1 \pmod n$ and consider the map $\varphi^* : \mathfrak{S}_n \to \mathfrak{S}_n$ defined by $\varphi^*(\sigma) = \sigma \circ \varphi$. 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 $0 \le k \le n-1$ are all torically equivalent.

We define a *linear extension*[1] of $[\alpha]_\sim$ to be a permutation σ such that there exists an acyclic orientation $\widehat{\alpha} \in [\alpha]_\sim$ with $\sigma \in \mathcal{L}(\widehat{\alpha})$. Letting $\mathcal{L}([\alpha]_\sim)$ denote the set of linear extensions of $[\alpha]_\sim$, we have

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

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. We can perform a *double flip* to an acyclic orientation α by choosing a source u and a sink v in α that are not adjacent to each other and then simultaneously flipping both of them. We say two acyclic orientations $\alpha, \alpha' \in \mathsf{Acyc}(G)$ are *double-flip equivalent*, denoted $\alpha \approx \alpha'$, if α' can be obtained from α via a sequence of double flips (and this, too, is an equivalence relation). Let $[\alpha]_\approx$ denote the equivalence class in $\mathsf{Acyc}(G)/\approx$ that contains α. Note that every equivalence class in $\mathsf{Acyc}/\sim$ is a union of equivalence classes in $\mathsf{Acyc}/\approx$. A *linear extension* of a double-flip equivalence class $[\alpha]_\approx$ is a permutation σ such that $\sigma \in \mathcal{L}(\widehat{\alpha})$ for some $\widehat{\alpha} \in [\alpha]_\approx$. Letting $\mathcal{L}([\alpha]_\approx)$ denote the set of linear extensions of $[\alpha]_\approx$, we have

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

The following theorem states that the connected components of $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ are parameterized by the double-flip equivalence classes of acyclic orientations of $\overline{Y}$ and that

[1] Note that our notion of a linear extension of a toric acyclic orientation differs from the definition of a "toric total order" in [6]. Indeed, that article defines a toric total order of $[\alpha]_\sim$ to be a cyclic equivalence class $\{(\varphi^*)^k(\sigma) : 0 \le k \le n-1\}$ such that σ is (using our definition) a linear extension of $[\alpha]_\sim$.

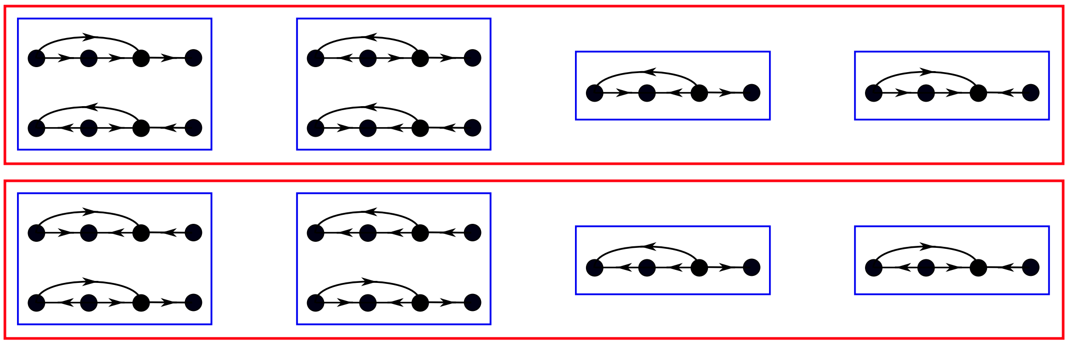

Figure 1: Each red box encompasses a toric acyclic orientation. Each blue box encompasses a double-flip equivalence class. As predicted by Theorem 3.3, each toric acyclic orientation is a union of 4 double-flip equivalence classes.

the vertices within a specific connected component are given by the linear extensions of the corresponding double-flip equivalence class. (Compare with Theorem 2.1.)

Theorem 3.1 ([5]). *Let Y be a graph with vertex set $[n]$. For each $[\alpha]_\approx \in \mathsf{Acyc}(\overline{Y})/\approx$, choose a linear extension $\sigma_{[\alpha]_\approx} \in \mathcal{L}([\alpha]_\approx)$, and let $H_{[\alpha]_\approx}$ be the connected component of $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ 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,*

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

The proof of Theorem 3.1 in [5] builds on the characterization of $\mathsf{FS}(\mathsf{Path}_n, Y)$ from Theorem 2.1. Since Cycle_n is obtained from Path_n by adding the edge $\{n, 1\}$, we see that the vertex set of each connected component of $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ is a union of the vertex sets of some connected components of $\mathsf{FS}(\mathsf{Path}_n, Y)$. Heuristically, two connected components of $\mathsf{FS}(\mathsf{Path}_n, Y)$ "merge" in $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ if they contain bijections that are related by a friendly swap across the edge $\{n, 1\}$. On the level of acyclic orientations, this added flexibility is precisely what the double-flip equivalence relation encodes.

Theorem 3.1 describes the connected components of $\mathsf{FS}(\mathsf{Cycle}_n, Y)$, but it is possible to say even more. Namely, we will obtain a description of these connected components that relies on only the flip equivalence relation $\sim$, not the double-flip equivalence relation $\approx$. The strategy of the proof is to begin with the special case in which $\overline{Y}$ is connected; here, each toric acyclic orientation of $\overline{Y}$ corresponds to n pairwise isomorphic connected components of $\mathsf{FS}(\mathsf{Cycle}_n, Y)$. For the general case where $\overline{Y}$ is not necessarily connected, we look at the local behavior for the connected components of $\overline{Y}$ separately and then stitch this information together using appropriate cyclic actions (whence the gcd). In what follows, if C is a subgraph of a friends-and-strangers graph $\mathsf{FS}(\mathsf{Cycle}_n, Y)$, we let $\varphi^*(C)$ denote the induced subgraph of $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ on the vertex set $\varphi^*(V(C))$.

Theorem 3.2 ([5]). *Let Y be a graph on the vertex set $[n]$. Let $n_1, \ldots, n_r$ denote the sizes of the connected components of $\overline{Y}$, and let $\nu = \gcd(n_1, \ldots, n_r)$. For each toric acyclic orientation $[\alpha]_\sim \in \mathsf{Acyc}(\overline{Y})/\sim$, choose a linear extension $\sigma_{[\alpha]_\sim}$ of $[\alpha]_\sim$, and let $J_{[\alpha]_\sim}$ be the connected component of $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ containing $\sigma_{[\alpha]_\sim}$. The graphs $J_{[\alpha]_\sim}, \varphi^*(J_{[\alpha]_\sim}), \ldots, (\varphi^*)^{\nu-1}(J_{[\alpha]_\sim})$ are distinct, pairwise isomorphic connected components of $\mathsf{FS}(\mathsf{Cycle}_n, Y)$. Moreover,*

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

The following result, which is of independent interest and not obvious *a priori*, provides a relationship between the double-flip equivalence classes of acyclic orientations of a graph G and the toric acyclic orientations of G. It is surprising that the proof of this fact passes through the analysis of the graph $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ with $Y = \overline{G}$.

Theorem 3.3 ([5]). *Let G be a graph with connected components of sizes $n_1, \ldots, n_r$, and let $\nu = \gcd(n_1, \ldots, n_r)$. Each toric acyclic orientation of G is the union of ν double-flip equivalence classes of $\mathsf{Acyc}(G)$.*

It is known that the number of toric acyclic orientations of a graph G is $T_G(1, 0)$, where T_G is the Tutte polynomial of G [6]. Thus, Theorems 3.1 and 3.3 tell us that if Y is a graph on $n \geq 3$ vertices such that the connected components of $\overline{Y}$ have sizes $n_1, \ldots, n_r$, then $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ has exactly $T_{\overline{Y}}(1, 0)\nu$ connected components, where $\nu = \gcd(n_1, \ldots, n_r)$. In particular, this enumeration yields a complete characterization of the graphs Y for which $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ is connected.

Corollary 3.4 ([5]). *Let Y be a graph with $n \geq 3$ vertices. Then $\mathsf{FS}(\mathsf{Cycle}_n, Y)$ is connected if and only if $\overline{Y}$ is a forest consisting of trees $\mathscr{T}_1, \ldots, \mathscr{T}_r$ such that $\gcd(|V(\mathscr{T}_1)|, \ldots, |V(\mathscr{T}_r)|) = 1$.*

4 General conditions for connectivity

The article [5] develops some necessary conditions and some sufficient conditions for the graph $\mathsf{FS}(X, Y)$ to be connected.

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 $\widetilde{X}$ such that:

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

We obtain a sufficient condition for $\mathsf{FS}(X, Y)$ to be connected in the case where X is a prolongation of a known small graph. Recall that a *hereditary class* is a collection of (isomorphism types of) graphs that is closed under taking induced subgraphs.

Theorem 4.1 ([5]). *Let $\mathcal{H}$ be a hereditary class. Let X be a graph on n_0 vertices with a Hamiltonian path, and suppose that $\mathsf{FS}(X, Y)$ is connected for every $Y \in \mathcal{H}$ on n_0 vertices. If $\widetilde{X}$ is a prolongation of X with n vertices, then $\mathsf{FS}(\widetilde{X}, \widetilde{Y})$ is connected for every $\widetilde{Y} \in \mathcal{H}$ on n vertices.*

We now mention a fairly general necessary condition for $\mathsf{FS}(X, Y)$ to be connected.

Theorem 4.2 ([5]). *Let X and Y be graphs on n vertices. Suppose $x_1 \cdots x_d$ ($d \geq 1$) is a path in X, where x_1 and x_d are cut vertices and each of $x_2, \ldots, x_{d-1}$ has degree exactly 2. If the minimum degree of Y is smaller than or equal to d, then $\mathsf{FS}(X, Y)$ is disconnected.*

These two theorems have several applications to understanding $\mathsf{FS}(X, Y)$ for particular choices of X, some of which are stated in [5]. One example is the lollipop graph $\mathsf{Lollipop}_{3,n-3}$, which is obtained from Path_n by adding the edge $\{n-2, n\}$; in this case, we have that $\mathsf{FS}(\mathsf{Lollipop}_{n-3,3}, Y)$ is connected if and only if the minimum degree of Y is at least $n-2$. It turns out that the same characterization holds when $X = \mathsf{D}_n$ is the Dynkin diagram of type D_n, which is obtained from Path_{n-1} by adding the additional vertex n and the additional edge $\{n-2, n\}$; here, we find that $\mathsf{FS}(\mathsf{D}_n, Y)$ is connected if and only if the minimum degree of Y is at least $n-2$.

5 Typical and extremal aspects

We next turn to the question of connectivity for $\mathsf{FS}(X, Y)$ when X and Y are random graphs. Our main result is the following. Recall that an event occurs *with high probability* if its probability of occurring tends to 1 as the sizes of the graphs involved tend to ∞.

Theorem 5.1 ([1]). *Fix $\varepsilon > 0$. Let X and Y be independently-chosen Erdős-Rényi random graphs on n vertices with edge probability $p = p(n)$. If $p \leq (2^{-1/2} - \varepsilon)n^{-1/2}$, then $\mathsf{FS}(X, Y)$ is disconnected with high probability. If $p \geq (\exp(2(\log n)^{2/3}))n^{-1/2}$, then $\mathsf{FS}(X, Y)$ is connected with high probability.*

We remark that the first statement in this theorem follows from a well-known result about edge-disjoint packings of graphs, which in this case correspond to isolated vertices in $\mathsf{FS}(X, Y)$; it seems that, as in the usual case of a binomial random graph, this local obstruction to connectedness tells essentially the whole story.

The proof of the second statement relies on the notion of an exchangeable pair of vertices, which generalizes the notion of an (X, Y)-friendly swap. Let X and Y be n-vertex graphs, and fix a bijection $\sigma \in V(\mathsf{FS}(X, Y))$. Let u and v be distinct vertices of Y, and write $u' = \sigma^{-1}(u)$ and $v' = \sigma^{-1}(v)$. Let $\sigma \circ (u' \; v')$ be the bijection that sends

u' to v, sends v' to u, and sends x to x for all $x \in V(X) \setminus \{u', v'\}$. We say that u and v are (X, Y)-*exchangeable from* σ if σ and $\sigma \circ (u' \; v')$ are in the same connected component of $\mathsf{FS}(X, Y)$. In other words, u and v are exchangeable from σ if there is a sequence of (X, Y)-friendly swaps that, when applied to σ, has the overall effect of swapping u and v (even if this swap is not itself (X, Y)-friendly). It turns out that having u and v be exchangeable for many bijections σ can be worth as much as having the edge $\{u, v\}$ in Y, as is illustrated by the following straightforward lemma.

Lemma 5.2 ([1]). *Let X and Y be n-vertex graphs, and suppose that X is connected. Suppose that for all distinct vertices $u, v \in V(Y)$ and every bijection σ satisfying $\{\sigma^{-1}(u), \sigma^{-1}(v)\} \in E(X)$, the vertices u and v are (X, Y)-exchangeable from σ. Then $\mathsf{FS}(X, Y)$ is connected.*

The other important idea in the proof of Theorem 5.1 concerns embedding small graphs in large random graphs. Let m be a positive integer, and let G and H be two graphs on the vertex set $[m]$. Let X and Y be n-vertex graphs, and let $\sigma : V(X) \to V(Y)$ be a bijection. Let $V_1, \ldots, V_m$ be a list of m pairwise disjoint sets of vertices of Y. We say that the pair of graphs (G, H) is *embeddable in (X, Y) with respect to the sets $V_1, \ldots, V_m$ and the bijection σ* if there exist vertices $v_i \in V_i$ for all $i \in [m]$ such that for all $i, j \in [m]$, we have

$$\{i, j\} \in E(H) \implies \{v_i, v_j\} \in E(Y) \quad \text{and} \quad \{i, j\} \in E(G) \implies \{\sigma^{-1}(v_i), \sigma^{-1}(v_j)\} \in E(X).$$

It is desirable to know if (G, H) is embeddable in (X, Y) with respect to the sets $V_1, \ldots, V_m$ and the bijection σ for all choices of σ and all reasonably large disjoint subsets $V_1, \ldots, V_m$. A technical argument using Janson's Inequalities shows that when X and Y are large random graphs and G and H are fixed small graphs, this is the case with high probability. (See [1] for a precise statement.)

The proof of the second statement in Theorem 5.1 proceeds by combining these two main ideas. We first find embeddings for particular specially-constructed pairs of sparse graphs (G, H) for which we know that certain pairs of vertices are exchangeable (and $|V(G)| = |V(H)| \approx (\log n)^{2/3}$ grows slowly with n). By choosing our embeddings carefully, we show that with high probability any two vertices u and v in Y are exchangeable; the argument involves lifting u and v to (G, H), finding a sequence of swaps that exchanges them there, and then bringing this sequence back down to (X, Y). Finally, we use Lemma 5.2 to conclude that $\mathsf{FS}(X, Y)$ is connected.

We also mention an analogous result for the case where X and Y are random edge-subgraphs of the complete bipartite graph $K_{r,r}$. It is not hard to show that if X and Y are both bipartite graphs on $n \geq 3$ vertices, then $\mathsf{FS}(X, Y)$ has multiple connected components. Hence, it makes sense to investigate when $\mathsf{FS}(X, Y)$ has the minimum possible number of connected components, namely, 2.

Theorem 5.3 ([1]). *Fix $\varepsilon > 0$. Let X and Y be independently-chosen Erdős-Rényi random edge-subgraphs of $K_{r,r}$ with edge probability $p = p(r)$. If $p \leq (1 - \varepsilon) r^{-1/2}$, then $\mathsf{FS}(X, Y)$ has more* *than* 2 *connected components with high probability. If* $p \geq (5(\log r)^{1/10})r^{-3/10}$*, then* $\mathsf{FS}(X,Y)$ *has exactly* 2 *connected components with high probability.*

Finally, we address the question of determining minimum degree conditions on X and Y that guarantee the connectedness of $\mathsf{FS}(X,Y)$.

Theorem 5.4 ([1]). *For each* $n \geq 1$*, let* d_n *denote the smallest nonnegative integer such that whenever* X *and* Y *are* n*-vertex graphs each with minimum degree at least* d_n*, the graph* $\mathsf{FS}(X,Y)$ *is connected. We have* $d_n \geq 3n/5 - 2$*. If* $n \geq 16$*, then* $d_n \leq 9n/14 + 2$*.*

We also establish a bipartite analogue which, in this case, is very nearly exact.

Theorem 5.5 ([1]). *For each* $r \geq 2$*, let* $d_{r,r}$ *be the smallest nonnegative integer such that whenever* X *and* Y *are edge-subgraphs of* $K_{r,r}$ *each with minimum degree at least* $d_{r,r}$*, the graph* $\mathsf{FS}(X,Y)$ *has exactly* 2 *connected components. We have* $\lceil (3r+1)/4 \rceil \leq d_{r,r} \leq \lceil (3r+2)/4 \rceil$*.*

For each theorem, the lower bound comes from an explicit construction (see [1] for more details), and the upper bound relies on delicate technical arguments in which lack of structure in some part of (X,Y) forces more rigid structure elsewhere. The notion of exchangeability again plays a central role.

6 Conclusion and future directions

The articles [1] and [5] contain several conjectures and open problems. Some have already been resolved, and we hope that others will lead to further developments in the study of friends-and-strangers graphs. We list a few of these questions here.

In Sections 2 and 3, we analyzed the connected components of $\mathsf{FS}(X,Y)$ where X is a path or a cycle, and Wilson [8] has analyzed the case where X is a star; other choices may also be fruitful for future inquiry.

Problem 6.1. *Characterize the connected components of* $\mathsf{FS}(X,Y)$ *for more fixed choices of* X*.*

It is easy to see that $\mathsf{FS}(K_n,Y)$ is connected if and only if Y is connected, but the bipartite setting seems to introduce more subtleties.

Problem 6.2. *Characterize the edge-subgraphs* Y *of* $K_{r,r}$ *such that* $\mathsf{FS}(K_{r,r},Y)$ *has exactly* 2 *connected components.*

The following conjecture, about determining the graphs X such that $\mathsf{FS}(X,Y)$ is connected for all "reasonable" choices of Y, has recently been proven by Ryan Jeong (personal communication).

Conjecture 6.3 ([5]). *Let* X *be a graph on* $n \geq 3$ *vertices. If* $\mathsf{FS}(X,\mathsf{Cycle}_n)$ *is connected, then* $\mathsf{FS}(X,Y)$ *is connected for every biconnected graph* Y*.*

The following conjecture pertains to the bounds in Section 5.

Conjecture 6.4 ([1])**.** *There is an absolute constant $C > 0$ such that if X and Y are independently chosen Erdős-Rényi random graphs on n vertices (respectively, Erdős-Rényi random edge-subgraphs of $K_{r,r}$) with edge probability $p \geq Cn^{-1/2}$, then $\mathsf{FS}(X,Y)$ is connected (respectively, has exactly 2 connected components) with high probability.*

In an exciting recent development, Bangachev [2] has studied asymmetric minimum degree conditions on X and Y that guarantee the connectedness of $\mathsf{FS}(X,Y)$. One consequence of his results is the following improvement of our Theorem 5.4.

Theorem 6.5 ([2])**.** *We have (in the notation of Theorem 5.4) that $d_n \leq \lceil 3n/5 \rceil$.*

Besides the number of connected components, another natural parameter of friends-and-strangers graphs that could be worth considering is the diameter.

Question 6.6 ([1])**.** *Does there exist an absolute constant $C > 0$ such that for all n-vertex graphs X and Y, every connected component of $\mathsf{FS}(X,Y)$ has diameter at most n^C?*

Finally, we mention that it could be fruitful to study random walks on friends-and-strangers graphs; indeed, this corresponds to friends and strangers randomly walking on graphs. Random walks on $\mathsf{FS}(X,K_n)$ correspond to the interchange process on X as discussed, for example, in [3].

References

[1] N. Alon, C. Defant, and N. Kravitz. "Typical and extremal aspects of friends-and-strangers graphs". 2020. arXiv:2009.07840.

[2] K. Bangachev. "On the asymmetric generalizations of two extremal questions on friends-and-strangers graphs". 2021. arXiv:2107.06789.

[3] P. Caputo, T. M. Liggett, and T. Richthammer. "Proof of Aldous' spectral gap conjecture". *J. Amer. Math. Soc.* **23** (2010), pp. 831–851.

[4] P. Cartier and D. Foata. *Problèmes combinatoires de commutation et réarrangements*. Berlin–New York: Springer–Verlag, 1969.

[5] C. Defant and N. Kravitz. "Friends and strangers walking on graphs". *Comb. Theory* (2021). To appear. arXiv:2009.05040.

[6] M. Develin, M. Macauley, and V. Reiner. "Toric partial orders". *Trans. Amer. Math. Soc.* **368** (2016), pp. 2263–2287.

[7] R. P. Stanley. "An equivalence relation on the symmetric group and multiplicity-free flag *h*-vectors". *J. Comb.* **3** (2012), pp. 277–298.

[8] R. M. Wilson. "Graph puzzles, homotopy, and the alternating group". *J. Combin. Theory, Ser. B* **16** (1974), pp. 86–96.