

CONNECTEDNESS AND CYCLE SPACES OF FRIENDS-AND-STRANGERS...

COLIN DEFANT

Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
e-mail: colindefant@gmail.com

DAVID DONG

Eastside Preparatory School, Kirkland, WA 98004, USA
e-mail: etsptq@gmail.com

ALAN LEE

Henry M. Gunn High School, Palo Alto, CA 94306, USA
e-mail: alandongjinlee@gmail.com

AND

MICHELLE WEI

The Harker School, San Jose, CA 95124, USA
e-mail: michelle.wei89@gmail.com

Abstract

If $X = (V(X), E(X))$ and $Y = (V(Y), E(Y))$ are n -vertex graphs, then their *friends-and-strangers graph* $\text{FS}(X, Y)$ is the graph whose vertices are the bijections from $V(X)$ to $V(Y)$ in which two bijections σ and σ' are adjacent if and only if there is an edge $\{a, b\} \in E(X)$ such that $\{\sigma(a), \sigma(b)\} \in E(Y)$ and $\sigma' = \sigma \circ (a\ b)$, where $(a\ b)$ is the permutation of $V(X)$ that swaps a and b . We prove general theorems that provide necessary and/or sufficient conditions for $\text{FS}(X, Y)$ to be connected. As a corollary, we obtain a complete characterization of the graphs Y such that $\text{FS}(\text{Dand}_{k,n}, Y)$ is connected, where $\text{Dand}_{k,n}$ is a dandelion graph; this substantially generalizes a theorem of the first author and Kravitz in the case $k = 3$. For specific choices of Y , we characterize the spider graphs X such that $\text{FS}(X, Y)$ is

connected. In a different vein, we study the cycle spaces of friends-and-strangers graphs. Naatz proved that if X is a path graph, then the cycle space of $\text{FS}(X, Y)$ is spanned by 4-cycles and 6-cycles; we show that the same statement holds when X is a cycle and Y has domination number at least 3. When X is a cycle and Y has domination number at least 2, our proof sheds light on how walks in $\text{FS}(X, Y)$ behave under certain Coxeter moves.

Keywords: friends-and-strangers graph, spider, dandelion, cycle space, Coxeter move.

2020 Mathematics Subject Classification: 05C40, 05C70.

1. INTRODUCTION

Flip graphs are graphs that encode when combinatorial objects are related by small changes called *flips*. In recent years, these graphs have received a great deal of attention in combinatorics, geometry, and computer science. For example, some of the most well-studied flip graphs are the 1-skeleta of polytopes such as the permutahedron, the associahedron, and the cyclohedron. In [5], the first author and Kravitz introduced a broad family of flip graphs called friends-and-strangers graphs.

Suppose we have simple graphs $X = (V(X), E(X))$ and $Y = (V(Y), E(Y))$ such that $|V(X)| = |V(Y)| = n$. We imagine that the vertices of X are chairs and that the vertices of Y are people; two people are adjacent in Y if and only if they are friends with each other (otherwise, they are strangers). The *friends-and-strangers graph* of X and Y , denoted $\text{FS}(X, Y)$, is a graph whose vertices are the bijections from $V(X)$ to $V(Y)$; one can think of such a bijection as an arrangement of people sitting in chairs. If we are given such an arrangement, then we allow two people to swap places with each other if they are friends with each other and they are sitting in adjacent chairs. Such a swap is called an (X, Y) -friendly swap. The edges of $\text{FS}(X, Y)$ correspond precisely to (X, Y) -friendly swaps. More formally, two bijections $\sigma, \sigma' : V(X) \rightarrow V(Y)$ are adjacent in $\text{FS}(X, Y)$ if and only if there exists an edge $\{a, b\} \in E(X)$ such that $\{\sigma(a), \sigma(b)\} \in E(Y)$, $\sigma(a) = \sigma'(b)$, $\sigma(b) = \sigma'(a)$, and $\sigma(c) = \sigma'(c)$ for all $c \in V(X) \setminus \{a, b\}$. For example, suppose

$$X = Y = \overset{1}{\bullet} - \overset{2}{\bullet} - \overset{3}{\bullet}.$$

Since $V(X) = V(Y) = \{1, 2, 3\}$, we can represent each bijection $\sigma : V(X) \rightarrow V(Y)$ as a permutation $\sigma(1)\sigma(2)\sigma(3)$ in one-line notation. Then

$$\text{FS}(X, Y) = \begin{array}{ccc} 213 & 123 & 132 \\ \bullet & - & \bullet & - & \bullet \\ 312 & 321 & 231 \end{array}.$$

Friends-and-strangers graphs generalize several previously-studied notions. For example, when $Y = K_n$ is the complete graph with vertex set $[n] := \{1, \dots, n\}$ and X is another graph with vertex set $[n]$, the friends-and-strangers graph $\text{FS}(X, Y)$ is the Cayley graph of the symmetric group \mathfrak{S}_n generated by the collection of transpositions corresponding to the edges of X . The famous 15-puzzle is equivalent to analyzing $\text{FS}(\text{Star}_{16}, \text{Grid}_{4 \times 4})$, where Star_n is the star graph with n vertices and $\text{Grid}_{4 \times 4}$ is the 4×4 grid graph. Generalizing the 15-puzzle, Wilson [12] characterized the graphs Y such that $\text{FS}(\text{Star}_n, Y)$ is connected (see Theorem 12). Stanley [10] studied the connected components of $\text{FS}(\text{Path}_n, \text{Path}_n)$, where Path_n is the path graph with n vertices.

It is very natural to ask about the connected components of a friends-and-strangers graph; indeed, if we view the vertices of $\text{FS}(X, Y)$ as arrangements of people in chairs, then two such arrangements are in the same connected component of $\text{FS}(X, Y)$ if and only if one can be obtained from the other via a sequence of (X, Y) -friendly swaps. As mentioned above, Wilson studied this problem when X is a star graph. Several papers have continued this line of work when one of X or Y is fixed to be a specific type of graph [5, 7], when X and Y are Erdős–Rényi random graphs [1, 11], or when X and Y satisfy certain minimum-degree conditions [1, 3]. Jeong recently studied the girths and diameters of friends-and-strangers graphs [8]. The first author has also related friends-and-strangers graphs of the form $\text{FS}(\text{Cycle}_n, Y)$, where Cycle_n is the cycle with n vertices, to a dynamical system called toric promotion [4].

In [5], the first author and Kravitz found sufficient conditions for $\text{FS}(X, Y)$ to be connected under the hypothesis that X has a Hamiltonian path. One of our primary goals in this paper is to prove the following sufficient condition for $\text{FS}(X, Y)$ to be connected; this condition is very general and is quite different from the ones established in [5].

Theorem 1. *Let X and Y be connected n -vertex graphs such that X has maximum degree $k \geq 2$. If every k -vertex induced subgraph of Y is connected and there exists a k -vertex induced subgraph Y_0 of Y such that $\text{FS}(\text{Star}_k, Y_0)$ is connected, then $\text{FS}(X, Y)$ is connected.*

In practice, Wilson’s theorem, which we have recorded as Theorem 12 below, makes it easy to check whether or not $\text{FS}(\text{Star}_k, Y_0)$ is connected for a given graph Y_0 . In fact, Wilson’s result will allow us to prove the following, which immediately lets us ignore one of the hypotheses of Theorem 1 when $n \geq 2k - 1$.

Theorem 2. *Let Y be an n -vertex graph such that every k -vertex induced subgraph of Y is connected. If $n \geq 2k - 1$, then there exists a k -vertex induced subgraph Y_0 of Y such that $\text{FS}(\text{Star}_k, Y_0)$ is connected.*

The hypothesis that $n \geq 2k - 1$ in the previous theorem is necessary; indeed,

the conclusion of the theorem is false (by Theorem 12 below) if Y is the complete bipartite graph $K_{k-1,k-1}$.

A *spider* is a connected graph that has a vertex c such that deleting c results in a disjoint union of paths. The vertex c is called the *center* of the spider, and the paths that result from deleting c are called the *legs* of the spider (the center is unique if there are at least three legs). The number of vertices in a leg is called its *length*. We write $\text{Spider}(\lambda_1, \dots, \lambda_k)$ for the spider with legs of lengths $\lambda_1, \dots, \lambda_k$; note that the number of vertices in this graph is $\lambda_1 + \dots + \lambda_k + 1$. The star graph Star_n is the spider with $n - 1$ legs, all of which have length 1.

While Theorem 1 provides a sufficient condition for the connectedness of $\text{FS}(X, Y)$, the next theorem provides a necessary condition when X is a spider.

Theorem 3. *Let $\lambda_1 \geq \dots \geq \lambda_k$ be positive integers, and let $n = \lambda_1 + \dots + \lambda_k + 1$. Let Y be an n -vertex graph. If there exists a disconnected induced subgraph of Y with $n - \lambda_1$ vertices, then $\text{FS}(\text{Spider}(\lambda_1, \dots, \lambda_k), Y)$ is disconnected.*

The *dandelion* graph $\text{Dand}_{k,n}$ is the spider graph with $k - 1$ legs of length 1 and 1 leg of length $n - k$. In [5], it was shown that the graph $\text{FS}(\text{Dand}_{2,n}, Y)$ is connected if and only if Y is the complete graph K_n (note that $\text{Dand}_{2,n} = \text{Path}_n$) and that for $n \geq 5$, the graph $\text{FS}(\text{Dand}_{3,n}, Y)$ is connected if and only if the minimum degree of Y is at least $n - 2$. The following corollary, which follows from Theorems 1, 2, and 3, substantially generalizes these results.

Corollary 4. *If $n \geq 2k - 1$, then $\text{FS}(\text{Dand}_{k,n}, Y)$ is connected if and only if every induced subgraph of Y with k vertices is connected.*

Corollary 4 is noteworthy because it is rare to find families of graphs X such that we can completely characterize all graphs Y such that $\text{FS}(X, Y)$ is connected. The next few theorems provide further characterizations of connectedness, though they put restrictions on both X and Y . In what follows, we denote \overline{G} to be the *complement* of G , which has the same vertex set as G and satisfies $\{u, v\} \in E(\overline{G})$ if and only if $\{u, v\} \notin E(G)$.

Theorem 5. *Let $\lambda_1 \geq \dots \geq \lambda_k$ be positive integers such that $\lambda_1 + \dots + \lambda_k + 1 = n \geq 4$. The friends-and-strangers graph $\text{FS}(\text{Spider}(\lambda_1, \dots, \lambda_k), \overline{\text{Cycle}_n})$ is connected if and only if $(\lambda_1, \dots, \lambda_k)$ is not of the form $(\lambda_1, 1, 1)$ and is not in the following list:*

$$(1, 1, 1, 1), \quad (2, 2, 1), \quad (2, 2, 2), \quad (3, 2, 1), \quad (3, 3, 1), \quad (4, 2, 1), \quad (5, 2, 1).$$

For $n \geq 4$, define the *fruit graph* Cycle_n^\perp to be the graph obtained from a cycle of size $n - 1$ by adding an extra vertex and a single edge connecting that vertex to one of the vertices in the cycle. More precisely, Cycle_n^\perp has vertex set $[n]$ and edge set

$$\{1, n - 1\} \cup \{1, n\} \cup \{\{i, i + 1\} : i \in [n - 2]\}.$$

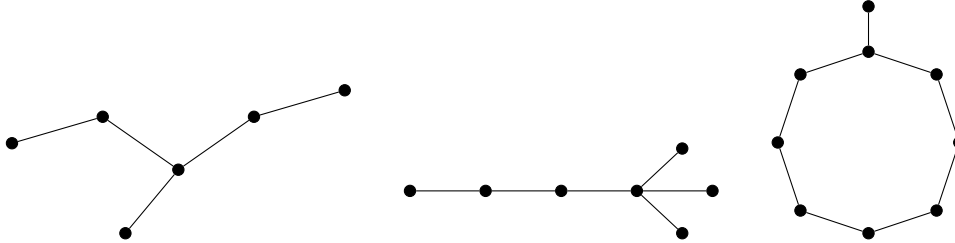


Figure 1. From left to right: $\text{Spider}(2, 2, 1)$, $\text{Dand}_{4,7}$, Cycle_9^\perp .

Figure 1 shows examples of a spider, a dandelion, and a fruit graph. The following result characterizes when the friends-and-strangers graph of a spider and the complement of a fruit graph is connected.

Theorem 6. *Let $\lambda_1 \geq \dots \geq \lambda_k$ be positive integers such that $k \geq 3$ and $\lambda_1 + \dots + \lambda_k + 1 = n$. Then $\text{FS}(\text{Spider}(\lambda_1, \dots, \lambda_k), \text{Cycle}_n^\perp)$ is disconnected if and only if $(\lambda_1, \dots, \lambda_k)$ is of one of the following forms:*

$$(\lambda_1, 1, 1, 1), \quad (\lambda_1, \lambda_2, 1), \quad (2, 2, 2).$$

The next theorem guarantees the connectedness of $\text{FS}(X, Y)$ whenever the minimum degree of Y is large and X is a connected graph that contains some small spider.

Theorem 7. *Let X and Y be n -vertex graphs such that Y has minimum degree at least $n - 3$. The friends-and-strangers graph $\text{FS}(X, Y)$ is connected if X is connected and contains a (not necessarily induced) subgraph isomorphic to at least one of the following:*

$$\begin{aligned} &\text{Star}_7, \quad \text{Spider}(2, 1, 1, 1, 1), \quad \text{Spider}(2, 2, 1, 1), \\ &\text{Spider}(3, 3, 2), \quad \text{Spider}(4, 2, 2), \quad \text{Spider}(4, 3, 1). \end{aligned}$$

All of our results mentioned so far have been concerned with whether or not a friends-and-stranger graph is connected. In a different vein, it is also natural to study cycles in friends-and-strangers graphs; see [2, 6, 9] for previous work on cycles in other flip graphs. Note that a cycle in a friends-and-strangers graph represents a nontrivial way that we can perform a sequence of friendly swaps that returns us to the arrangement of people with which we started.

A graph is called *even-degree* if each of its vertices has even degree. An *edge-subgraph* of a graph G is a subgraph of G that has the same vertex set as G . Given edge-subgraphs H and H' of G , let $H \triangle H'$ be the edge-subgraph of G whose edge set is the symmetric difference of the edge sets of H and H' . Note that if H and H' are even-degree, then so is $H \triangle H'$. The *cycle space* of G is the

set of all even-degree edge-subgraphs; it is a vector space over the 2-element field \mathbb{F}_2 in which the addition operation is given by the symmetric difference Δ . We can view a cycle in G as an edge-subgraph in which all vertices not in the cycle are isolated vertices. It is well known that the cycle space of a (finite) graph is spanned by its cycles. Naatz [9] proved that if Y is any n -vertex graph, then the cycle space of $\text{FS}(\text{Path}_n, Y)$ is spanned by 4-cycles and 6-cycles (he stated this in the case when Y is the incomparability graph of an n -element poset, but his methods apply more generally for any Y).

In this article, we will study cycle spaces of friends-and-strangers graphs of the form $\text{FS}(\text{Cycle}_n, Y)$. The requisite analysis ends up being more complicated than what Naatz used to study $\text{FS}(\text{Path}_n, Y)$, but we will still obtain analogues of several of his results.

A *dominating set* of a graph G is a subset D of the vertex set of G such that every vertex in G is either in D or is adjacent to a vertex in D . The minimum size of a dominating set of G is called the *domination number* of G . In the following theorem, we must impose the additional (fairly mild) condition that the domination number of Y is at least 3.

Theorem 8. *Let Y be an n -vertex graph with domination number at least 3. The cycle space of $\text{FS}(\text{Cycle}_n, Y)$ is spanned by 4-cycles and 6-cycles. If Y is triangle-free, then the cycle space of $\text{FS}(\text{Cycle}_n, Y)$ is spanned by 4-cycles.*

The proof of Theorem 8 proceeds by first establishing a more general result (Theorem 20) about when walks in $\text{FS}(\text{Cycle}_n, Y)$ can be obtained from one another via a sequence of Coxeter moves, which are essentially the relations defining the affine symmetric group as a Coxeter group.

As an example of Theorem 8, let Y be the dandelion $\text{Dand}_{3,8}$, which has domination number 3 and is triangle-free. Figure 2 shows one connected component of $\text{FS}(\text{Cycle}_8, \text{Dand}_{3,8})$; upon inspection, we see that the cycle space of this connected component is generated by 4-cycles, as predicted by Theorem 8.

Remark 9. The hypothesis that the domination number of Y is at least 3 in Theorem 8 is necessary; indeed the theorem fails for $\text{FS}(\text{Cycle}_4, \text{Cycle}_4)$.

Remark 10. It would be interesting to understand the cycle spaces of other types of friends-and-strangers graphs besides those of the form $\text{FS}(\text{Path}_n, Y)$ or $\text{FS}(\text{Cycle}_n, Y)$.

The plan for the paper is as follows. In Section 2, we establish some notation and terminology, and we recall some previous results about friends-and-strangers graphs. In Section 3, we prove Theorems 1, 2, and 3 and use them to deduce Corollary 4. In Section 4, we establish Theorems 5, 6, and 7. To do so, we first prove a general lemma (Lemma 16) that allows us to build connected friends-and-strangers graphs from smaller ones. Section 5 studies Coxeter moves and

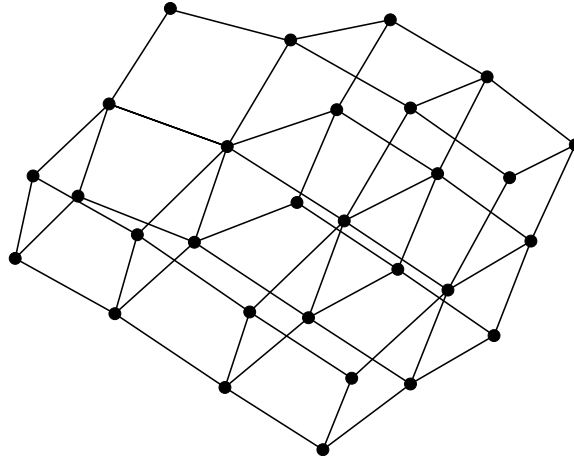


Figure 2. One connected component of $\text{FS}(\text{Cycle}_8, \text{Dand}_{3,8})$.

cycle spaces of friends-and-strangers graphs of the form $\text{FS}(\text{Cycle}_n, Y)$; it is in this section that we prove Theorem 8.

2. PRELIMINARIES

The purpose of this section is to fix some terminology and notation and to discuss previous results that we will need in future sections.

2.1. Transpositions and walks

For any set S and elements $s, s' \in S$, we write $(s \ s')$ for the transposition—which we view as a bijection from S to itself—that swaps s and s' and fixes all elements of $S \setminus \{s, s'\}$. This is useful for notating friendly swaps in the friends-and-strangers graph $\text{FS}(X, Y)$. Indeed, suppose σ is a vertex in $\text{FS}(X, Y)$ and $\{x, x'\} \in E(X)$. Let $y = \sigma(x)$ and $y' = \sigma(x')$. If $\{y, y'\} \in E(Y)$, then we can perform an (X, Y) -friendly swap across the edge $\{x, x'\}$ in order to obtain a new permutation $\sigma' = \sigma \circ (x \ x') = (y \ y') \circ \sigma$.

A *walk* in a graph G is a sequence of vertices such that any two consecutive vertices in the sequence are adjacent in G . Equivalently, we can think of a walk as a starting vertex together with a sequence of edges such that any two consecutive edges share a vertex. Each edge e in the friends-and-strangers graph $\text{FS}(X, Y)$ corresponds to an (X, Y) -friendly swap. It is convenient to label such an edge with the pair of people (i.e., vertices of Y) who performed the swap. More precisely, if $e = \{\sigma, \sigma'\}$, then we define the *edge label* $\psi(e)$ to be the unique pair

$\{y, y'\}$ such that $\sigma' = (y \ y') \circ \sigma$. For ease of notation, we will often omit the set braces and write yy' (or equivalently, $y'y$) for this edge label $\{y, y'\}$. Note that we can specify a walk in $\text{FS}(X, Y)$ by saying its starting vertex together with its *label sequence*, which is just the sequence of edge labels of the edges used in the walk.

2.2. Bipartite graphs

The following proposition from [5] tells us that the friends-and-strangers graph of two bipartite graphs is disconnected.

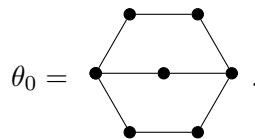
Proposition 11 [5]. Suppose X and Y are bipartite graphs with $V(X) = V(Y) = [n]$. Let $\{A_X, B_X\}$ be a bipartition of X , and let $\{A_Y, B_Y\}$ be a bipartition of Y . For $\sigma \in V(\text{FS}(X, Y))$, define $p(\sigma) \in \{0, 1\}$ by

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

where $\text{sgn}(\sigma)$ is the sign of σ , viewed as a permutation in the symmetric group S_n . If $\tau, \tau' \in V(\text{FS}(X, Y))$ are such that $p(\tau) \neq p(\tau')$, then τ and τ' are in different connected components of $\text{FS}(X, Y)$.

2.3. Wilson's theorem

Given a graph G and a vertex $v \in V(G)$, we write $G \setminus v$ for the graph obtained from G by deleting v (and all edges incident to v). We say v is a *cut vertex* of G if G is connected and $G \setminus v$ is disconnected. We say G is *biconnected* if it is connected and has no cut vertices. Wilson noted that if $n \geq 3$ and Y is an n -vertex graph that is not biconnected, then $\text{FS}(\text{Star}_n, Y)$ is disconnected. Thus, he focused his attention on friends-and-strangers graphs of stars and biconnected graphs. The statement of his main theorem involves the exceptional graph



Theorem 12 [12]. 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 with $n!/2$ vertices. The graph $\text{FS}(\text{Star}_7, \theta_0)$ has exactly 6 connected components.

Remark 13. It is straightforward to describe the connected components of $\text{FS}(\text{Star}_n, \text{Cycle}_n)$. For each vertex σ of $\text{FS}(\text{Star}_n, \text{Cycle}_n)$, we can read off the leaves of Star_n in the clockwise order that their images under σ appear around

Cycle_n . This defines a cyclic ordering of the set of leaves of Star_n . It is straightforward to see that two permutations are in the same connected component of $\text{FS}(\text{Star}_n, \text{Cycle}_n)$ if and only if they induce the same cyclic ordering of the leaves of Star_n .

3. GENERAL THEOREMS ON SPIDERS

Our goal in this section is to prove Theorems 1, 2, and 3, which together imply Corollary 4. We begin with Theorem 2 since its proof does not require any further preliminary results.

Proof of Theorem 2. Suppose $n \geq 2k - 1$, and let Y be an n -vertex graph such that every k -vertex induced subgraph of Y is connected. We will use induction on ℓ to show that for each $\ell \in \{3, \dots, k\}$, there exists a biconnected ℓ -vertex induced subgraph $Y^{(\ell)}$ of Y that contains a triangle. The proof will then follow by setting $Y_0 = Y^{(k)}$. Indeed, if $\ell = 3$, then $\text{FS}(\text{Star}_3, Y_0) = \text{FS}(\text{Star}_3, K_3)$ is connected, and if $\ell \geq 4$, then Theorem 12 guarantees that $\text{FS}(\text{Star}_k, Y_0)$ is connected (since $Y^{(k)}$ is biconnected, not bipartite, not a cycle, and not θ_0).

Because every k -vertex induced subgraph of Y is connected, we see that every vertex in Y has degree at least $n - k + 1$. The number of edges in Y is half the sum of the degrees of the vertices in Y , which is at least $n(n - k + 1)/2 > n^2/4$. Therefore, it follows from Turán's theorem that Y contains a triangle $Y^{(3)}$. This completes the base case of our induction.

Now suppose $\ell \in \{4, \dots, k\}$, and assume inductively that there exists a biconnected $(\ell - 1)$ -vertex induced subgraph $Y^{(\ell-1)}$ of Y that contains a triangle. We claim that the number of edges between vertices in $V(Y^{(\ell-1)})$ and vertices in $V(Y) \setminus V(Y^{(\ell-1)})$ is at least $n - \ell + 2$. Since $|V(Y) \setminus V(Y^{(\ell-1)})| = n - \ell + 1$, it will then follow from the pigeonhole principle that there exists a vertex $v \in V(Y) \setminus V(Y^{(\ell-1)})$ that has at least two neighbors in $Y^{(\ell-1)}$. We can then simply set $Y^{(\ell)}$ to be the induced subgraph of Y whose vertex set is $V(Y^{(\ell-1)}) \cup \{v\}$; it is easy to see that this induced subgraph is biconnected and contains a triangle (the same triangle as in $Y^{(\ell-1)}$).

To prove that claim, we use the fact, which we mentioned before, that each vertex in Y has degree at least $n - k + 1$. Since $Y^{(\ell-1)}$ has $\ell - 1$ vertices, this implies that each vertex in $Y^{(\ell-1)}$ has at least $n - k + 1 - (\ell - 2) = n - k - \ell + 3$ neighbors in $V(Y) \setminus V(Y^{(\ell-1)})$. It follows that there are at least $(\ell - 1)(n - k - \ell + 3)$ edges between vertices in $V(Y^{(\ell-1)})$ and vertices in $V(Y) \setminus V(Y^{(\ell-1)})$. Thus, to prove the claim, we must show that $(\ell - 1)(n - k - \ell + 3) \geq n - \ell + 2$. This inequality rearranges to $p(\ell) \geq 0$, where $p(x) = -x^2 + (n - k + 5)x - (2n - k + 5)$. We will show that $p(3) \geq 0$ and $p(k) \geq 0$; since $p(x)$ is a quadratic polynomial and $3 \leq \ell \leq k$, this will imply that $p(\ell) \geq 0$. First, we have $p(3) = n - 2k + 1 \geq 0$

by hypothesis. On the other hand, $p(k) = q(k)$, where $q(x)$ is the quadratic polynomial $-2x^2 + (n+6)x - (2n+5)$. To see that $q(k) \geq 0$, we use the fact that $3 \leq k \leq (n+1)/2$ and compute that $q(3) = n-5 \geq 0$ and $q((n+1)/2) = (n-5)/2 \geq 0$. ■

We now move on to proving Theorems 1, and 3, for which we need the following lemmas.

Lemma 14. *Fix an integer $n \geq 3$, and let Y be a biconnected graph with n vertices. Let σ be a vertex of $\text{FS}(\text{Star}_n, Y)$, and fix $x \in V(\text{Star}_n)$ and $y \in V(Y)$. There exists a vertex σ' in the same connected component of $\text{FS}(\text{Star}_n, Y)$ as σ such that $\sigma'(x) = y$.*

Proof. If Y is a cycle, then the desired result follows easily from Remark 13. If $n = 7$ and Y is the exceptional graph θ_0 shown in Section 2.3, then we can check by hand that the desired result still holds. Now suppose Y is not a cycle or θ_0 . If Y is not bipartite, then Theorem 12 tells us that $\text{FS}(X, Y)$ is connected, so the desired result is obvious.

Now suppose Y is bipartite. Since Y is biconnected, we must have $n \geq 4$. We may assume $V(\text{Star}_n) = V(Y) = [n]$; let $\{A_Y, B_Y\}$ be a bipartition of Y . We assume that n is the center of Star_n so that $\{[n-1], \{n\}\}$ is a bipartition of Star_n . For each $\mu \in V(\text{FS}(\text{Star}_n, Y))$, let $p(\mu) = |\mu([n-1]) \cap A_Y| + \frac{\text{sgn}(\mu)+1}{2} \pmod{2}$, as in Proposition 11. Choose $u, v \in [n-1] \setminus \{x\}$. Let τ be a vertex in $\text{FS}(X, Y)$ such that $\tau(x) = y$, and let $\tau' = \tau \circ (u \ v)$. Then $\tau'(x) = y$. We have $p(\tau) \neq p(\tau')$, so it follows from Proposition 11 that τ and τ' are in different connected components of $\text{FS}(\text{Star}_n, Y)$. Theorem 12 tells us that $\text{FS}(\text{Star}_n, Y)$ has exactly 2 connected components, so we can take σ' to be whichever of the vertices τ or τ' is in the same connected component as σ . ■

Lemma 15. *Let X and Y be connected n -vertex graphs such that X has maximum degree $k \geq 2$. Suppose every induced subgraph of Y with k vertices is connected. Let σ be a vertex of $\text{FS}(X, Y)$, and fix $x \in V(X)$ and $y \in V(Y)$. There exists a vertex σ' in the same connected component of $\text{FS}(X, Y)$ as σ such that $\sigma'(x) = y$.*

Proof. If T is a spanning tree of X with maximum degree k (it is straightforward to see that such a spanning tree exists), then $\text{FS}(T, Y)$ is a subgraph of $\text{FS}(X, Y)$. Therefore, it suffices to prove the lemma when X is a tree; we assume that this is the case in what follows. Note that the result is trivial if $\sigma(x) = y$. Since Y is connected, it suffices to prove the result when $\{\sigma(x), y\}$ is an edge in Y ; we assume that this is the case in what follows. We proceed by induction on n .

Let c be a vertex of X of degree k , and let $N[c]$ denote the closed neighborhood of c (i.e., the set of vertices that are adjacent to or equal to c). If X is a

star, then $n = k + 1$, so the hypothesis that every k -vertex induced subgraph of Y is connected is equivalent to the fact that Y is biconnected. Thus, the desired result follows immediately from Lemma 14. Note that this handles the base case $n = 3$. In what follows, we may assume X is not a star. In particular, X has a leaf $u \notin N[c]$. We consider three cases.

Case 1. Suppose X has a leaf u such that $u \notin N[c]$, $u \neq x$, and $u \neq \sigma^{-1}(y)$. Let $X' = X \setminus u$ and $Y' = Y \setminus \sigma(u)$. Let $\sigma|_{X'}$ be the restriction of σ to $V(X')$. Since u is a leaf of X and $N[c] \subseteq V(X')$, the graph X' is a tree with $n - 1$ vertices and maximum degree k . Furthermore, every k -vertex induced subgraph of Y' is connected. Therefore, it follows by induction that there exists a sequence of (X', Y') -friendly swaps that transforms $\sigma|_{X'}$ into a vertex μ of $\text{FS}(X', Y')$ such that $\mu(x) = y$. Let σ' be the unique vertex in $\text{FS}(X, Y)$ such that $\sigma'(u) = \sigma(u)$ and $\sigma'(v) = \mu(v)$ for all $v \in V(X')$. Then $\sigma'(x) = y$. The sequence of (X', Y') -friendly swaps transforming $\sigma|_{X'}$ into μ can be interpreted as a sequence of (X, Y) -friendly swaps that transforms σ into σ' . Thus, σ' is in the same connected component of $\text{FS}(X, Y)$ as σ .

Case 2. Suppose that x is a leaf of X and $x \notin N[c]$. Let x' be the neighbor of x in X . Let $X' = X \setminus x$ and $Y' = Y \setminus \sigma(x)$. Let $\sigma|_{X'}$ be the restriction of σ to $V(X')$. The graph X' is a tree with $n - 1$ vertices and maximum degree k , and every k -vertex induced subgraph of Y' is connected. It follows by induction (and the assumption $\sigma(x) \neq y$) that there exists a sequence of (X', Y') -friendly swaps that transforms $\sigma|_{X'}$ into a vertex μ of $\text{FS}(X', Y')$ such that $\mu(x') = y$. Let τ be the unique vertex in $\text{FS}(X, Y)$ such that $\tau(x) = \sigma(x)$ and $\tau(v) = \mu(v)$ for all $v \in V(X')$. In particular, $\tau(x') = y$. The sequence of (X', Y') -friendly swaps transforming $\sigma|_{X'}$ into μ can be interpreted as a sequence of (X, Y) -friendly swaps that transforms σ into τ . Thus, τ is in the same connected component of $\text{FS}(X, Y)$ as σ . Now let $\sigma' = \tau \circ (x \ x')$. Since $\{x, x'\} \in E(X)$ and $\{\tau(x), \tau(x')\} = \{\sigma(x), y\} \in E(Y)$, the vertices τ and σ' are adjacent in $\text{FS}(X, Y)$. Then σ' is in the same connected component as σ and satisfies $\sigma'(x) = y$.

Case 3. Suppose that there is a leaf u of X such that $u \notin N[c]$ and $\sigma(u) = y$. We have assumed that $\sigma(x) \neq y$, so $u \neq x$. Let y' be a neighbor of y in Y . Since $\{\sigma(u), y'\}$ is an edge in Y , we can repeat the argument in *Case 2* with x replaced by u and y replaced by y' ; this allows us to deduce that there is a vertex $\tilde{\sigma}$ of $\text{FS}(X, Y)$ in the same connected component as σ such that $\tilde{\sigma}(u) = y'$. But now observe that u is a leaf of X such that $u \notin N[c]$, $u \neq x$, and $u \neq \sigma^{-1}(y)$; this means that we can repeat the argument from *Case 1* with σ replaced by $\tilde{\sigma}$ to see that there is a vertex σ' in the same connected component of $\text{FS}(X, Y)$ as $\tilde{\sigma}$ such that $\sigma'(x) = y$. But then σ' is also in the same connected component of $\text{FS}(X, Y)$ as σ . ■

We can now complete the proofs of Theorems 1 and 3.

Proof of Theorem 1. Let n, k, X, Y, Y_0 be as in the statement of the theorem; we want to prove that $\text{FS}(X, Y)$ is connected. If T is a spanning tree of X , then $\text{FS}(T, Y)$ is a subgraph of $\text{FS}(X, Y)$. Therefore, we may assume in what follows that X is a tree. We may also assume that $n \geq 4$ since the case when $n = 3$ can be checked by hand.

First, suppose $n = k + 1$. In this case, X is isomorphic to Star_n . The hypothesis that every k -vertex induced subgraph of Y is connected is equivalent to the fact that Y is biconnected. Since $\text{FS}(\text{Star}_k, Y_0)$ is connected, it follows from the discussion in Section 2.3 that the graph Y cannot be a cycle or the exception graph θ_0 ; indeed, if it were, then Y_0 would either not be biconnected or would be a cycle. By Theorem 12 (or Proposition 11), the hypothesis that $\text{FS}(\text{Star}_k, Y_0)$ is connected guarantees that Y_0 is not bipartite, which implies that Y is not bipartite. Hence, it follows from Theorem 12 that $\text{FS}(\text{Star}_n, Y)$ is connected, as desired.

We may now assume $n \geq k + 2$ and proceed by induction on n . Let c be a vertex of X of degree k , and let $N[c]$ be the closed neighborhood of c . There must be a leaf x of X such that $x \notin N[c]$. Fix some vertex $y \in V(Y) \setminus \{Y_0\}$. Fix a vertex τ of $\text{FS}(X, Y)$ such that $\tau(x) = y$. Choose some vertex σ of $\text{FS}(X, Y)$. We will show that σ is in the same connected component as τ ; as σ was arbitrary, this will prove that $\text{FS}(X, Y)$ is connected.

According to Lemma 15, there is a vertex σ' in the same connected component of $\text{FS}(X, Y)$ as σ such that $\sigma'(x) = y$. Let $X' = X \setminus x$ and $Y' = Y \setminus y$. Let $\sigma'|_{X'}$ and $\tau|_{X'}$ be the restrictions of σ' and τ , respectively, to $V(X')$. Note that X' is a tree with $n - 1$ vertices and maximum degree k . Furthermore, every k -vertex induced subgraph of Y' is connected, and Y' contains the induced subgraph Y_0 such that $\text{FS}(\text{Star}_k, Y)$ is connected. Thus, we can use induction to see that $\text{FS}(X', Y')$ is connected. This means that there is a sequence of (X', Y') -friendly swaps that transforms $\sigma'|_{X'}$ into $\tau|_{X'}$. This sequence of (X', Y') -friendly swaps can be interpreted as a sequence of (X, Y) -friendly swaps that transforms σ' into τ . This shows that σ' is in the same connected component of $\text{FS}(X, Y)$ as τ , so σ is also in the same connected component as τ . ■

Proof of Theorem 3. Let $\lambda_1 \geq \dots \geq \lambda_k$ be as in the statement of the theorem, and let $X = \text{Spider}(\lambda_1, \dots, \lambda_k)$. Let z be the leaf of X on the leg of length λ_1 . Let us define a partial order \preceq on $V(X)$ by declaring that $x \preceq x'$ if the unique path from z to x' contains x .

By hypothesis, there is a disconnected induced subgraph Y_0 of Y with $n - \lambda_1$ vertices. There exist nonempty subsets $A, B \subseteq V(Y_0)$ such that $A \cup B = V(Y_0)$ and $A \cap B = \emptyset$ and such that no vertex in A is adjacent in Y to any vertex in B . Let us say a vertex $\sigma \in \text{FS}(X, Y)$ is *special* if there exists $a_0 \in A$ such that $\sigma^{-1}(a_0) \preceq \sigma^{-1}(b)$ for all $b \in B$. The graph $\text{FS}(X, Y)$ has at least one special vertex and at least one non-special vertex. Therefore, in order to prove that

$\text{FS}(X, Y)$ is disconnected, it suffices to show that every connected component of $\text{FS}(X, Y)$ consists entirely of special vertices or entirely of non-special vertices.

Suppose instead that there is a connected component of $\text{FS}(X, Y)$ that contains both special and non-special vertices. Then this connected component must have a special vertex σ that is adjacent to a non-special vertex τ . Let $\{u, v\} \in E(Y)$ be the edge label of the edge $\{\sigma, \tau\} \in E(\text{FS}(X, Y))$; that is, $\tau = \sigma \circ (u \ v)$. Then $\{\sigma^{-1}(u), \sigma^{-1}(v)\} \in E(X)$. Because σ is special, there exists $a_0 \in A$ such that $\sigma^{-1}(a_0) \preceq \sigma^{-1}(b)$ for all $b \in B$. Using the definition of the partial order \preceq and the fact that X is a spider, it is now straightforward to see that $a_0 \in \{u, v\}$; without loss of generality, we may assume $a_0 = u$. Note that $\tau^{-1}(v) = \sigma^{-1}(a_0) \preceq \sigma^{-1}(b) = \tau^{-1}(b)$ for all $b \in B$; if v were in A , then this would contradict the fact that τ is not special. This shows that $v \notin A$, and we also know that $v \notin B$ since $\{u, v\} \in E(Y)$ and no vertex in A is adjacent to any vertex in B . Thus, $v \in V(Y) \setminus (A \cup B)$.

Let L be the set of vertices in X that are in the leg of length λ_1 containing z . Thus, $|L| = \lambda_1$. Because σ is special and τ is not, there exists $b_0 \in B$ such that $\sigma^{-1}(a_0) \preceq \sigma^{-1}(b_0)$ while $\tau^{-1}(a_0) \not\preceq \tau^{-1}(b_0)$. This forces $\sigma^{-1}(a_0)$ to be the center of the spider X , and it also forces $\sigma^{-1}(v)$ to be in $V(X) \setminus L$. Since $\sigma^{-1}(a_0) \preceq \sigma^{-1}(b)$ for all $b \in B$, we must have $\tau^{-1}(B) = \sigma^{-1}(B) \subseteq V(X) \setminus L$. If there were a vertex $a_1 \in \tau^{-1}(A) \cap L$, then we would have $\tau^{-1}(a_1) \preceq \tau^{-1}(b)$ for all $b \in B$, contradicting the fact that τ is not special. Hence, $\tau^{-1}(A) \subseteq V(X) \setminus L$. We also know that $\tau^{-1}(v) = \sigma^{-1}(a_0)$ is the center of X , so $\tau^{-1}(v) \in V(X) \setminus L$. Thus, $\tau^{-1}(A) \cup \tau^{-1}(B) \cup \{\tau^{-1}(v)\}$ is a set of size $n - \lambda_1 + 1$ that is contained in $V(X) \setminus L$. This is a contradiction because $|V(X) \setminus L| = n - \lambda_1$. ■

4. SPIDERS, CYCLES, AND FRUITS

The goal of this section is to prove Theorems 5, 6, and 7. The main tool for proving these theorems is the following lemma, which will allow us to build connected friends-and-strangers graphs from smaller ones. A family \mathcal{Y} of (isomorphism classes of) graphs is called *hereditary* if it is closed under taking induced subgraphs (i.e., every induced subgraph of a graph in \mathcal{Y} is also in \mathcal{Y}).

Lemma 16. *Let \mathcal{Y} be a hereditary family of connected graphs. Let X be a graph with n vertices such that $\text{FS}(X, Y)$ is connected for every n -vertex graph $Y \in \mathcal{Y}$. Let $x \in V(X)$, and let X' be the graph obtained from X by adding a new vertex x' together with the edge $\{x, x'\}$. Then $\text{FS}(X', Y')$ is connected for every $(n+1)$ -vertex graph $Y' \in \mathcal{Y}$.*

Proof. We will prove that if $\sigma, \sigma' : V(X') \rightarrow V(Y')$ are two vertices of $\text{FS}(X', Y')$ such that $\sigma(x')$ and $\sigma'(x')$ are adjacent in Y' , then σ and σ' are in the same

connected component of $\text{FS}(X', Y')$. Since Y' is connected, this will imply that any two vertices in $\text{FS}(X', Y')$ is connected.

Let $y = \sigma(x')$ and $y' = \sigma'(x')$, and suppose $\{y, y'\} \in E(Y')$. Let F_y be the subgraph of $\text{FS}(X', Y')$ induced by the set of vertices τ satisfying $\tau(x') = y$. Similarly, define $F_{y'}$ to be the subgraph of $\text{FS}(X', Y')$ induced by the set of vertices τ' satisfying $\tau'(x') = y'$. Let Y_y and $Y_{y'}$ be the induced subgraphs of Y' on the vertex sets $V(Y') \setminus \{y\}$ and $V(Y') \setminus \{y'\}$, respectively. It is clear that F_y and $F_{y'}$ are isomorphic to $\text{FS}(X, Y_y)$ and $\text{FS}(X, Y_{y'})$, respectively. Since \mathcal{Y} is hereditary, we know that Y_y and $Y_{y'}$ are both in \mathcal{Y} . Therefore, our hypothesis on X guarantees that both F_y and $F_{y'}$ are connected graphs.

Let μ be any vertex in $\text{FS}(X', Y')$ satisfying $\mu(x) = y'$ and $\mu(x') = y$, and let μ' be the vertex $\mu \circ (x \ x')$. Note that $\mu'(x') = y'$ and that μ and μ' are adjacent in $\text{FS}(X, Y)$. We have $\mu \in F_y$ and $\mu' \in F_{y'}$, so there are paths in $\text{FS}(X', Y')$ from σ to μ and from σ' to μ' . Therefore, σ and σ' are in the same connected component of $\text{FS}(X', Y')$, as desired. ■

Corollary 17. *Let X be a graph with $n \geq 5$ vertices such that $\text{FS}(X, \overline{\text{Cycle}_n})$ is connected. Let $x \in V(X)$, and let X' be the graph obtained from X by adding a new vertex x' together with the edge $\{x, x'\}$. Then $\text{FS}(X', \overline{\text{Cycle}_{n+1}})$ is connected.*

Proof. Let \mathcal{Y} be the smallest hereditary family of graphs that contains $\overline{\text{Cycle}_N}$ for all $N \geq 5$. In other words, \mathcal{Y} is the collection of graphs that can be realized as induced subgraphs of complements of cycles. It is straightforward to verify that a graph is in \mathcal{Y} if and only if its complement is a cycle or a disjoint union of paths. It follows that if Y is an n -vertex graph in \mathcal{Y} , then $\overline{\text{Cycle}_n}$ is a subgraph of Y . This shows that $\text{FS}(X, Y)$ is connected for every n -vertex graph Y in \mathcal{Y} , so it follows from Lemma 16 that $\text{FS}(X', Y')$ is connected for every $(n+1)$ -vertex graph Y' in \mathcal{Y} . In particular, $\text{FS}(X', \overline{\text{Cycle}_{n+1}})$ is connected. ■

Corollary 18. *Let X be a graph with $n \geq 5$ vertices such that $\text{FS}(X, \overline{\text{Cycle}_n})$ and $\text{FS}(X, \overline{\text{Cycle}_n^\perp})$ are both connected. Let $x \in V(X)$, and let X' be the graph obtained from X by adding a new vertex x' together with the edge $\{x, x'\}$. Then $\text{FS}(X', \overline{\text{Cycle}_{n+1}^\perp})$ is connected.*

Proof. Let \mathcal{Y} be the smallest hereditary family of graphs that contains $\overline{\text{Cycle}_N^\perp}$ for every $N \geq 5$. In other words, \mathcal{Y} is the collection of graphs that can be realized as induced subgraphs of complements of fruit graphs. It is straightforward to verify that every n -vertex graph in \mathcal{Y} contains either $\overline{\text{Cycle}_n}$ or $\overline{\text{Cycle}_n^\perp}$ as a subgraph. This shows that $\text{FS}(X, Y)$ is connected for every n -vertex graph Y in \mathcal{Y} , so it follows from Lemma 16 that $\text{FS}(X', Y')$ is connected for every $(n+1)$ -vertex graph Y' in \mathcal{Y} . In particular, $\text{FS}(X', \overline{\text{Cycle}_{n+1}^\perp})$ is connected. ■

Corollary 19. *Let X be a graph with $n \geq 4$ vertices such that $\text{FS}(X, Y)$ is connected for all n -vertex graphs Y with minimum degree at least $n - 3$. Let $x \in V(X)$, and let X' be the graph obtained from X by adding a new vertex x' together with the edge $\{x, x'\}$. If Y' is a graph with $n + 1$ vertices and minimum degree at least $n - 2$, then $\text{FS}(X', Y')$ is connected.*

Proof. This follows immediately from Lemma 16 by setting \mathcal{Y} to be the collection of graphs whose complements have maximum degree at most 2. ■

We can now proceed to the proofs of the main theorems of this section.

Proof of Theorem 5. Choose $\lambda_1 \geq \dots \geq \lambda_k \geq 1$ with $k \geq 3$ and $\lambda_1 \dots + \lambda_k + 1 = n$. We can easily check using a computer that $\text{FS}(\text{Spider}(\lambda_1, \dots, \lambda_k), \overline{\text{Cycle}_n})$ is disconnected if $(\lambda_1, \dots, \lambda_k)$ is one of the seven partitions listed in the statement of the theorem. Now suppose $(\lambda_1, \dots, \lambda_k) = (\lambda_1, 1, 1)$. We know by [5, Theorem 6.5] that $\text{FS}(\text{Spider}(\lambda_1, 1, 1), Y)$ is connected if and only if the minimum degree of Y is at least $n - 2$; since the minimum degree of $\overline{\text{Cycle}_n}$ is $n - 3$, $\text{FS}(\text{Spider}(\lambda_1, 1, 1), \overline{\text{Cycle}_n})$ is disconnected.

Now assume $(\lambda_1, \dots, \lambda_k)$ is not of the form $(\lambda_1, 1, 1)$ and is not one of the seven partitions listed in the statement of Theorem 5. It is straightforward to check that these assumptions guarantee that $\text{Spider}(\lambda_1, \dots, \lambda_k)$ contains an induced subgraph $\text{Spider}(\rho_1, \dots, \rho_r)$, where (ρ_1, \dots, ρ_r) is one of the following partitions:

$$(1, 1, 1, 1, 1), \quad (2, 1, 1, 1), \quad (6, 2, 1), \quad (4, 3, 1), \quad (3, 2, 2).$$

Let $n' = \rho_1 + \dots + \rho_r + 1$. We can verify by computer that

$$\text{FS}(\text{Spider}(\rho_1, \dots, \rho_r), \overline{\text{Cycle}_{n'}})$$

is connected. Now, $\text{Spider}(\lambda_1, \dots, \lambda_k)$ can be obtained from $\text{Spider}(\rho_1, \dots, \rho_r)$ by a sequence of operations, where each operation adds a new vertex to the graph and adds a new edge that has the new vertex as one of its endpoints. Therefore, it follows from Corollary 17 that $\text{FS}(\text{Spider}(\lambda_1, \dots, \lambda_k), \overline{\text{Cycle}_n})$ is connected. ■

Proof of Theorem 6. We can use a computer to check that

$$\begin{aligned} &\text{FS}(\text{Spider}(1, 1, 1, 1, 1), \overline{\text{Cycle}_6^\perp}), \\ &\text{FS}(\text{Spider}(2, 2, 1, 1), \overline{\text{Cycle}_7^\perp}), \\ &\text{FS}(\text{Spider}(2, 2, 3), \overline{\text{Cycle}_8^\perp}) \end{aligned}$$

are connected. It now follows easily from Corollary 18 that the graph

$$\text{FS}(\text{Spider}(\lambda_1, \dots, \lambda_k), \overline{\text{Cycle}_n^\perp})$$

is connected whenever $k \geq 5$ or $k = 4$ and $\lambda_2 \geq 2$ or $k = 3$ and $\lambda_1 \geq 3$ and $\lambda_3 \geq 2$.

With a computer, we can also verify that $\text{FS}(\text{Spider}(2, 2, 2), \overline{\text{Cycle}_7^\perp})$ has 12 connected components. Thus, to prove the reverse direction, we must show that $\text{FS}(\text{Spider}(\lambda_1, \dots, \lambda_k), \overline{\text{Cycle}_n^\perp})$ is disconnected if $k = 3$ and $\lambda_3 = 1$ or if $k = 4$ and $\lambda_2 = 1$. Since $V(\overline{\text{Cycle}_n^\perp}) = V(\text{Cycle}_n^\perp)$, we can consider the identity bijection $\text{id}: V(\text{Cycle}_n^\perp) \rightarrow V(\overline{\text{Cycle}_n^\perp})$; it is straightforward to see that this bijection is an isolated vertex in $\text{FS}(\text{Cycle}_n^\perp, \overline{\text{Cycle}_n^\perp})$. Hence, $\text{FS}(\text{Cycle}_n^\perp, \overline{\text{Cycle}_n^\perp})$ is disconnected. It is also straightforward to see that $\text{Spider}(\lambda_1, \lambda_2, 1)$ is a subgraph of Cycle_n^\perp , so this resolves the case where $k = 3$ and $\lambda_3 = 1$. Now suppose $k = 4$ and $\lambda_2 = 1$. Then $\text{Spider}(\lambda_1, 1, 1, 1)$ is the dandelion $\text{Dand}_{4,n}$, and $n - \lambda_1 = 4$. Since Cycle_n^\perp clearly contains a disconnected induced subgraph with 4 vertices, it follows from Theorem 3 that $\text{FS}(\text{Spider}(\lambda_1, 1, 1, 1), \overline{\text{Cycle}_n^\perp})$ is disconnected. ■

Proof of Theorem 7. Let G be one of the six graphs listed in the statement of Theorem 7, and let X be a connected graph that contains G as a subgraph. Let n and n' be the number of vertices of X and G , respectively. We can check by computer that $\text{FS}(G, H)$ is connected for every graph H with n' vertices and minimum degree at least $n' - 3$. Because X is connected, it is possible to obtain an n -vertex subgraph \tilde{X} of X by sequence of operations, where each operation adds a new vertex to the graph and adds a new edge that has the new vertex as one of its endpoints. It follows from Corollary 19 that if Y is an n -vertex graph with minimum degree at least $n - 3$, then $\text{FS}(\tilde{X}, Y)$ is connected. Since \tilde{X} is a subgraph of X with the same vertex set as X , the friends-and-strangers graph $\text{FS}(X, Y)$ must also be connected for every such Y . ■

5. COXETER MOVES AND THE CYCLE SPACE

Recall from Section 1 that the *cycle space* of a graph G is the vector space over \mathbb{F}_2 whose elements are the even-degree edge-subgraphs of G and whose addition is the symmetric difference operation Δ . This vector space is spanned by the cycles of G . The goal of this section is to prove Theorem 8, which provides smaller generating sets for the cycle spaces of graphs of the form $\text{FS}(\text{Cycle}_n, Y)$ when Y has domination number at least 3.

We provide a framework to study paths in $\text{FS}(\text{Cycle}_n, Y)$ that is inspired by Coxeter moves in the affine symmetric group. Under this framework, we also prove some results that apply when Y has domination number 2.

Recall the definitions of *walks*, *edge labels*, and *label sequences* from Section 2. We are going to describe certain moves that one can perform on a walk W to obtain a new walk with the same starting and ending vertices. The reader familiar

with the theory of Coxeter groups should recognize these operations as essentially Coxeter moves in the affine symmetric group, whose Coxeter graph is Cycle_n (though we will not need this formalism). If there are two consecutive identical edges in W , then we can simply delete them; we call this operation a *square deletion* (because it corresponds to the fact that the square of a simple reflection in the affine symmetric group is the identity). The opposite of a square deletion is a *square insertion*, which consists of inserting two consecutive identical edges into the walk so that the resulting sequence of edges is a walk. If there are two consecutive edges in W that do not share a vertex, then we can simply swap those two edges; we call this operation a *commutation move*. Finally, if W contains three consecutive edge labels of the form ab, ac, bc , then we can replace them with the edges bc, ac, ab ; we call this operation a *Yang–Baxter move*. Collectively, we refer to square deletions, square insertions, commutation moves, and Yang–Baxter moves as *Coxeter moves*.

One should think of a commutation move as replacing two edges in a 4-cycle in $\text{FS}(\text{Cycle}_n, Y)$ with the other two edges in the same 4-cycle. Similarly, one should think of a Yang–Baxter move as replacing three edges in a 6-cycle in $\text{FS}(\text{Cycle}_n, Y)$ with the other three edges in the same 6-cycle.

We will often specify a walk in $\text{FS}(\text{Cycle}_n, Y)$ by listing its label sequence. It is helpful to note that a Yang–Baxter move acts on the label sequence of a walk by simply reversing the order of three consecutive edge labels of the form ab, ac, bc .

We define an *anchored walk* to be a walk in $\text{FS}(\text{Cycle}_n, Y)$ whose label sequence is of the form $a_1b_1, a_2b_2, \dots, a_kb_k, a_1b_1$; the first and last edge labels in this walk (which are identical) are called the *anchors*. We say that such an anchored walk is *repetition-free* if the labels $a_1b_1, a_2b_2, \dots, a_kb_k$ are distinct. An anchored walk is *trivial* if its label sequence has only two (necessarily identical) labels. We can perform Coxeter moves on anchored walks in the same way that we performed them on ordinary walks; however, for anchored walks, we preserve the anchors and thus possibly shorten the walk. For example, suppose an anchored walk starts at a vertex σ and uses edges with labels ab, ac, bc, ab . The anchors are the first and last labels ab . We can perform a Yang–Baxter move on the first three edge labels to transform this walk into bc, ac, ab, ab . However, as an anchored walk, we would cut out the edge labels that are no longer between the two ab anchors. Thus, the new anchored walk would have label sequence ab, ab , so it would be trivial.

Let us say an anchored walk in $\text{FS}(\text{Cycle}_n, Y)$ is *complete* if its label sequence is of the form $ab, au_1, \dots, au_k, bu_{k+1}, \dots, bu_{n-2}, ab$, where $\{a, b, u_1, \dots, u_{n-2}\} = [n] = V(\text{Cycle}_n)$. Note that the existence of such a complete walk forces $\{a, b\}$ to be a dominating set of Y . We say an anchored walk W *reduces* to an anchored walk W' if there is a sequence of Coxeter moves that transforms the label sequence

of W into that of W' such that the number of square insertions that insert two copies of the label xy is at most the number of square deletions that delete two copies of the label xy .

Theorem 20. *Let Y be an n -vertex graph with domination number at least 2, and let W be a repetition-free anchored walk in $\text{FS}(\text{Cycle}_n, Y)$ with anchors ab . If there exists a vertex $z \in V(Y)$ such that the label sequence of W either contains both az and bz or contains neither az nor bz , then W reduces to a trivial anchored walk. Otherwise, W reduces to a complete anchored walk.*

In order to prove Theorem 20, it will be convenient to consider some particular subsequences of an anchored walk. Let Y be a graph with vertex set $[n]$. Suppose W is a repetition-free anchored walk in $\text{FS}(\text{Cycle}_n, Y)$ whose anchors are the label ab , where $a < b$. We define the *strong essential prefix* of W to be the longest initial subsequence of the label sequence of W of the form

$$ab, ax_1, \dots, ax_k, by_1, \dots, by_\ell$$

such that $a, b, x_1, \dots, x_k, y_1, \dots, y_\ell$ are all distinct. The *strong essential suffix* of W is the part of the label sequence of W not in the strong essential prefix. Similarly, we define the *weak essential prefix* of W to be the longest initial subsequence of the label sequence of W of the form $ab, ax_1, \dots, ax_k, by_1, \dots, by_\ell, c_1d_1, \dots, c_rd_r$ such that $a, b, x_1, \dots, x_k, y_1, \dots, y_\ell$ are all distinct and $c_1, \dots, c_r, d_1, \dots, d_r \in \{x_1, \dots, x_k, y_1, \dots, y_\ell\}$. The *weak essential suffix* is the part of the label sequence of W not in the weak essential prefix. Note that the strong essential prefix is a subsequence of the weak essential prefix and that the weak essential suffix is a subsequence of the strong essential suffix.

Lemma 21. *Let Y be a graph with vertex set $[n]$. Let W be a repetition-free anchored walk in $\text{FS}(\text{Cycle}_n, Y)$ with anchors ab , where $a < b$. If there exists $z \in V(Y) \setminus \{a, b\}$ such that neither az nor bz appear in the label sequence of W , then W reduces to a trivial anchored walk.*

Proof. We proceed by induction on the length of W . By definition, W is a sequence $\sigma_0, \dots, \sigma_p$ of vertices in $\text{FS}(\text{Cycle}_n, Y)$. One should think of each vertex σ_i as an arrangement of the vertices of Y on the vertices of Cycle_n . In the starting vertex σ_0 , the vertices a and b are sitting on adjacent vertices of Cycle_n , and they swap places when we move to σ_1 . At the end of the walk, we reach σ_p from σ_{p-1} by swapping a and b again. We have assumed that the label sequence does not contain az or bz . This readily implies that if the label sequence of W contains a label of the form ax with $x \neq b$, then it must also contain bx . Similarly, if the label sequence contains a label by such that $y \neq a$, then it also contains ay . This shows that it suffices to prove that W reduces to an anchored walk whose label

sequence does not contain az or bz and whose strong essential suffix is just the single label ab since such an anchored walk must be trivial. Let m be the length of the strong essential suffix of W ; we may assume $m \geq 2$ since otherwise we are done. By induction on m , it suffices to show that W reduces to an anchored walk W' whose label sequence does not contain az or bz and whose strong essential suffix has length at most $m - 1$.

Let $ab, ax_1, \dots, ax_k, by_1, \dots, by_\ell$ be the strong essential prefix of W . Let uv be the first label in the strong essential suffix of W . In other words, u and v are the two vertices of Y that take part in the friendly swap used to get from $\sigma_{k+\ell+1}$ to $\sigma_{k+\ell+2}$. We know that $uv \neq ab$ because $m \geq 2$. Let Q be the part of the strong essential suffix of W that comes after the label uv . We now consider three cases.

Case 1. Suppose $\{u, v\} \cap \{a, b, x_1, \dots, x_k, y_1, \dots, y_\ell\} = \emptyset$. In this case, we can repeatedly apply commutation moves so that each move swaps the label uv with the label to its left. After we perform $k + \ell + 1$ such moves, the label uv will leave the anchored walk, resulting in a new anchored walk W' whose strong essential suffix is contained in Q .

Case 2. Suppose $\{u, v\} \cap \{x_1, \dots, x_k, y_1, \dots, y_\ell\} = \emptyset$ and $\{u, v\} \cap \{a, b\} \neq \emptyset$. We will assume $u = a$ since the other case is similar. We can perform commutation moves to transform the label sequence of W into $ab, ax_1, \dots, ax_k, uv, by_1, \dots, by_\ell, Q$; this is the label sequence of the desired anchored walk W' (note that its strong essential suffix has length at most $m - 1$ because its strong essential prefix contains $ab, ax_1, \dots, ax_k, uv, by_1, \dots, by_\ell$).

Case 3. Suppose $\{u, v\} \cap \{x_1, \dots, x_k, y_1, \dots, y_\ell\} \neq \emptyset$. We will assume that $\{u, v\} \cap \{y_1, \dots, y_\ell\} \neq \emptyset$; the case when $\{u, v\} \cap \{x_1, \dots, x_k\} \neq \emptyset$ is similar. Let t be the largest index such that $y_t \in \{u, v\}$. Without loss of generality, we may assume $y_t = v$. The fact that W is repetition-free implies that $u \neq b$. Also, the maximality of t guarantees that $u \neq y_{t+1}$ (when $t < \ell$). Since $\sigma_{k+\ell+2}$ is obtained from $\sigma_{k+\ell+1}$ by performing a friendly swap involving the vertices u and $v = y_t$, we know that $\sigma_{k+\ell+1}^{-1}(u)$ must be one of the two vertices of Cycle_n adjacent to $\sigma_{k+\ell+1}^{-1}(y_t)$. One of these vertices is either $\sigma_{k+\ell+1}^{-1}(b)$ (if $t = \ell$) or $\sigma_{k+\ell+1}^{-1}(y_{t+1})$ (if $t < \ell$); since u is not b or y_{t+1} , it follows that $\sigma_{k+\ell+1}^{-1}(u)$ is the other of these two vertices.

Suppose first that $t \geq 2$. Then $u = y_{t-1}$. We can perform a sequence of commutation moves in order to move the label uv to the left until it is immediately to the right of by_t . We can then perform a Yang–Baxter move to transform the subsequence by_{t-1}, by_t, uv into uv, by_t, by_{t-1} . Finally, since the sets $\{u, v\}$ and $\{a, b, x_1, \dots, x_k, y_1, \dots, y_{t-2}\}$ are disjoint, we can apply further commutation moves in order to move uv to the left until it leaves the anchored walk. This results in the desired anchored walk W' .

Next, suppose $t = 1$ and $k = 0$. Then $u = a$. We can use commutation moves to transform the label sequence of W into $ab, by_1, uv, by_2, \dots, by_\ell, Q$. Since $uv = ay_1$, we can use a Yang–Baxter move to turn this sequence into $uv, by_1, ab, by_2, \dots, by_\ell, Q$. The label sequence of the new anchored walk W' does not include the labels uv and by_1 , so its strong essential suffix has length at most $m - 1$. Furthermore, the label sequence of W' does not contain az or bz .

Finally, suppose $t = 1$ and $k \geq 1$. Then $u = x_1$, so $uv = x_1y_1$. Let Q' be the sequence $ax_2, \dots, ax_k, by_2, \dots, by_\ell, Q$. Because ax_1 is in the label sequence of W , it follows from the first paragraph of this proof that bx_1 appears in Q ; hence, b and x_1 must be adjacent in Y . Furthermore, we know $\sigma_2^{-1}(b)$ is adjacent to $\sigma_2^{-1}(x_1)$. Then we can use a square insertion to change $ab, ax_1, by_1, x_1y_1, Q'$ into $ab, ax_1, bx_1, bx_1, by_1, x_1y_1, Q'$. Applying Yang–Baxter moves to the subsequences ab, ax_1, bx_1 and bx_1, by_1, x_1y_1 yields $bx_1, ax_1, ab, x_1y_1, by_1, bx_1, Q'$. Another commutation move then yields $bx_1, ax_1, x_1y_1, ab, by_1, bx_1, Q'$. This shows that we can apply Coxeter moves to transform W into an anchored walk whose label sequence is $ab, by_1, bx_1, ax_2, \dots, ax_k, by_2, \dots, by_\ell, Q$. Performing additional commutation moves yields an anchored walk W_0 whose label sequence is

$$ab, ax_2, \dots, ax_k, by_1, bx_1, by_2, \dots, by_\ell, Q.$$

Now, to meet the definition of what it means for an anchored walk to *reduce* to another, we must show that we can perform additional Coxeter moves to delete a pair of labels bx_1 . As mentioned before, the label bx_1 must appear in Q . Let W_1 be the anchored walk whose label sequence is the subsequence of W_0 starting and ending at the labels bx_1 . Note that since W_1 does not use the labels az or bz , it cannot contain the label x_1z . Because W_1 is repetition-free and does not use the labels x_1z or bz , we can apply induction on the length of the anchored walk to find that W_1 reduces to a trivial anchored walk. This means that we can perform a sequence of Coxeter moves to the label sequence of W_0 in order to move the two occurrences of the label bx_1 next to each other. We can then perform a square deletion to remove the two occurrences of bx_1 . Let W_2 be the resulting anchored walk with anchors ab (the same as the anchors of W_0). When we reduced W_1 to the trivial anchored walk, all labels that we added by square insertions were later deleted by square deletions (by the definition of *reducing*). This implies that W reduces to W_2 , that W_2 is repetition-free, and that W_2 does not use the labels az or bz . Since the length of W_2 is strictly less than that of W , we can use induction to see that W_2 , and hence also W , reduces to a trivial anchored walk (whose strong essential suffix has length 1). ■

Lemma 22. *Let W be a repetition-free anchored walk in $\text{FS}(\text{Cycle}_n, Y)$. Let*

$$ab, ax_1, \dots, ax_k, by_1, \dots, by_\ell, c_1d_1, \dots, c_rd_r$$

be the weak essential prefix of W , where

$$c_1, \dots, c_r, d_1, \dots, d_r \in \{x_1, \dots, x_k, y_1, \dots, y_\ell\}.$$

If $k \geq 1$ and the weak essential suffix of W begins with bu for some

$$u \in \{x_1, \dots, x_k\},$$

then W reduces to a trivial anchored walk. Similarly, if $\ell \geq 1$ and the weak essential suffix of W begins with au for some $u \in \{y_1, \dots, y_\ell\}$, then W reduces to a trivial anchored walk.

Proof. We prove only the first statement since the second is similar. Our strategy is to manipulate W through Coxeter moves and apply Lemma 21. Let W be the sequence of vertices $\sigma_0, \dots, \sigma_p$. First, we wish to apply Coxeter moves to move au left until it appears directly after the first anchor. This is already true if $u = x_1$. If $u = x_i$, where $i \geq 2$, then let Q be the subsequence of the label sequence of W consisting of all labels before ax_{i-1} , and let R be the subsequence consisting of all labels after au . Because W is repetition-free, the labels ux_1, \dots, ux_{i-1} must all occur in the label sequence between the labels au and bu . This implies that $\{u, x_j\} \in E(Y)$ for all $j \in [i-1]$. Thus, to move the label ax_i before ax_{i-1} , we can use a square insertion to turn the subsequence ax_{i-1}, au into $x_{i-1}u, x_{i-1}u, ax_{i-1}, au$, followed by a Yang–Baxter move on the last three labels, which changes this into $x_{i-1}u, au, ax_{i-1}, x_{i-1}u$. The first occurrence of $x_{i-1}u$ is disjoint from all labels in Q , so we remove it from W via commutation moves. The sequence now begins with $Q, au, ax_{i-1}, x_{i-1}u$. We repeat this process until au appears directly after ab , obtaining a new anchored walk W' with the label sequence $ab, au, ax_1, ux_1, ax_2, ux_2, \dots, ax_{i-1}, ux_{i-1}, R$. Now, we must delete all edges that we added with square insertions. For $j \in [i-1]$, let W_j be the anchored walk with the two occurrences of ux_j in W' as its anchors (the label ux_j occurs exactly once in R for each j). The only repeated labels in W' besides the anchors are ux_j for $j \in [i-1]$, so there are no repeated labels in W_{i-1} besides the anchors. Furthermore, the labels au and ax_{i-1} are outside of W_{i-1} . Thus, by Lemma 21, W_{i-1} reduces to a trivial anchored walk, and we can remove its anchors with a square deletion. Now that the labels ux_{i-1} have been deleted, the same argument applies to the anchored walk W_{i-2} , and so on; in this way, for each $j \in [i-1]$, we can use a square deletion to remove a pair of edge labels ux_j, ux_j . Thus, W reduces to a new anchored walk W'' that begins with ab, au and contains bu .

Having au directly after ab allows us to use another square insertion to change the subsequence ab, au into bu, bu, ab, au . Applying a Yang–Baxter move transforms this into bu, au, ab, bu . Now, because W' initially contained the label bu , the new anchored walk now contains two instances of bu ; we have created an

anchored walk B within W'' with anchors bu . Both au and ab appear outside of B , so by Lemma 21, B reduces to a trivial anchored walk, and its anchors can be removed with a square deletion. The labels bu and au do not appear in the resulting anchored walk. Thus, W'' reduces to a trivial anchored walk, so the initial walk W also reduces to a trivial anchored walk. ■

Now, we are ready to prove the main theorem.

Proof of Theorem 20. If there is some z such that neither az nor bz appear in W , then W reduces to a trivial anchored walk by Lemma 21. For all other W , we use induction show that W reduces to an anchored walk that either satisfies the hypothesis of Lemma 22 or has only the ending anchor ab as its weak essential suffix. Let $ab, ax_1, \dots, ax_k, by_1, \dots, by_\ell, c_1d_1, \dots, c_rd_r$ be the weak essential prefix of W , where $c_1, \dots, c_r, d_1, \dots, d_r \in \{x_1, \dots, x_k, y_1, \dots, y_\ell\}$. Let m be the length of the weak essential suffix of W . If $m = 1$, then we are done. If $m \geq 2$, we consider the first label uv of the weak essential suffix. If $k \geq 1$ and this label is bu for some $u \in \{x_1, \dots, x_k\}$ or $\ell \geq 1$ and this label is au for some $u \in \{y_1, \dots, y_\ell\}$, then we are done; W reduces to a trivial anchored walk by Lemma 22. Otherwise, we consider three cases in a similar way as in the proof of Lemma 21.

Case 1. Suppose $\{u, v\} \cap \{a, b, x_1, \dots, x_k, y_1, \dots, y_\ell\} = \emptyset$. We can apply commutation moves to remove uv from the anchored walk in a similar way as in the proof of Lemma 21.

Case 2. Suppose $\{u, v\} \cap \{x_1, \dots, x_k, y_1, \dots, y_\ell\} = \emptyset$ and $\{u, v\} \cap \{a, b\} \neq \emptyset$. We can again handle this case in a similar way as in Lemma 21 by performing commutation moves to move uv directly after ax_k .

Case 3. Suppose $\{u, v\} \cap \{x_1, \dots, x_k, y_1, \dots, y_\ell\} \neq \emptyset$. Using the same logic as in the proof of Lemma 21, we see that the only possibilities for uv are ay_1 if $k = 0$ and $\ell \geq 1$, bx_1 if $k \geq 1$ and $\ell = 0$, $x_t x_{t-1}$ for some $t \geq 2$, $y_t y_{t-1}$ for some $t \geq 2$, or $x_1 y_1$. The cases where $uv = ay_1$ or bx_1 are both handled by Lemma 22. By the definition of a weak essential prefix, uv cannot be $x_t x_{t-1}$, $y_t y_{t-1}$, or $x_1 y_1$.

If, throughout the inductive process, W never satisfies the condition in the statement of Lemma 22, then the weak essential suffix of W is just the ending anchor ab . We can then apply commutation moves to move each label $c_i d_i$ so that they appear after ab . Now, the strong essential suffix of the resulting anchored walk W' is ab , so W' must be trivial or complete. Furthermore, we know that for each $x \notin \{a, b\}$, either ax or bx (not both) must occur in W' , which forces W' to be a complete anchored walk. ■

Now, we can apply Theorem 20 to obtain results for when the domination number of Y is at least 3. Recall that a *geodesic path* in a graph G is a path that has minimum length among all paths with the same endpoints.

Proposition 23. *Let Y be an n -vertex graph with domination number at least 3. Let W be a geodesic path in $\text{FS}(\text{Cycle}_n, Y)$. Then no two edges used in W have the same label.*

Proof. Suppose instead that there is a label that appears twice in the label sequence of W . Then we can find a subsequence of the label sequence of W that is the label sequence of a repetition-free anchored walk W' . Let ab be the anchor label of W' . Since $\{a, b\}$ is not a dominating set of Y , there must be a vertex z of Y that is not adjacent to a or b in Y . Then the labels az and bz cannot appear in the label sequence of W' , so it follows from Theorem 20 that W' reduces to a trivial anchored walk. Once we perform a sequence of Coxeter moves to reduce W' to a trivial anchored walk, we can delete the two consecutive labels ab . We can think of performing all of the Coxeter moves on the walk W (which is not necessarily an anchored walk). These moves transform W into a new walk with the same endpoints as W that uses fewer edges than W , contradicting the assumption that W is a geodesic path. ■

Proposition 24. *Suppose Y is an n -vertex graph with domination number at least 3, and let C be a cycle in $\text{FS}(\text{Cycle}_n, Y)$. For every edge e in C , there exists an edge $e' \neq e$ in C such that $\psi(e) = \psi(e')$.*

Proof. Let $e = \{\tau, \tau'\}$, and let $\psi(e) = \{a, b\} \in E(Y)$. Let W be the unique walk in $\text{FS}(\text{Cycle}_n, Y)$ that starts at τ , ends at τ' , and uses each edge in C other than e exactly once. We want to show that one of the edges in W has edge label $\{a, b\}$. It is convenient to once again view the vertices of Y as people who are walking on Cycle_n . The arrangement τ' is obtained from τ by swapping the people a and b , who are standing on adjacent vertices of Cycle_n . Without loss of generality, say $a = \tau(1) = \tau'(2)$ and $b = \tau(2) = \tau'(1)$. As we traverse the walk W , we can keep track of the positions of a and b on the cycle. It is straightforward to see that if none of the edges in this walk have label $\{a, b\}$, then each person in $V(Y) \setminus \{a, b\}$ must have swapped with either a or b . However, since $\{a, b\} \in E(Y)$, this would imply that $\{a, b\}$ is a dominating set of Y , contradicting our assumption that Y has domination number at least 3. ■

Every friends-and-strangers graph is bipartite since its vertices can be viewed as permutations in a symmetric group and performing a friendly swap corresponds to multiplying by a transposition (which changes the sign of the permutation). This tells us that every cycle C in a friends-and-strangers graph has an even number of edges, so it makes sense to talk about pairs of opposite edges and pairs of opposite vertices of C . An *isometric cycle* of a graph G is a subgraph H of G that is a cycle and has the property that for all vertices u and v of H , the distance between u and v in H is the same as the distance between u and v in

G. The following proposition is an analogue of one that initially appears in [9] for $\text{FS}(\text{Path}_n, Y)$.

Proposition 25. *Let Y be an n -vertex graph with domination number at least 3. Let C be an isometric cycle in $\text{FS}(\text{Cycle}_n, Y)$. Two edges in C have the same edge label if and only if they are opposite edges of C .*

Proof. Let $2m$ be the number of edges in C . Let e and f be opposite edges in C . Proposition 24 tells us that there is some edge $e' \neq e$ in C with the same edge label as e . Let W_1, \dots, W_m be the paths within C that have m edges and that use the edge e (note that there are exactly m such paths). Because C is an isometric cycle of $\text{FS}(\text{Cycle}_n, Y)$, each path W_i is a geodesic path in $\text{FS}(\text{Cycle}_n, Y)$. It follows from Proposition 23 that e' is not an edge in any of the paths W_1, \dots, W_m . However, the only edge in C that is not in any of the paths W_1, \dots, W_m is f , so $e' = f$. This shows that e and f have the same edge label and that no other edges in C have the same edge label as e . ■

Lemma 26. *Let Y be an n -vertex graph with domination number at least 3. Let C be a cycle of length at least 8 in $\text{FS}(\text{Cycle}_n, Y)$. Either there are at least two consecutive edges in C that are contained in a common 4-cycle, or there are at least three consecutive edges in C that are contained in a common 6-cycle. Moreover, if Y is triangle-free, then there are at least two consecutive edges in C that are contained in a common 4-cycle.*

Proof. Our proof follows that of [9, Lemma 4.8]. If there are two consecutive edges in C that have disjoint edge labels, then they must belong to a 4-cycle. Now assume that the edge labels of any two consecutive edges in C have exactly one vertex of Y in common. We will show that Y contains a triangle and that there are at least three consecutive edges in C that are contained in a common 6-cycle.

Let e_1, \dots, e_k be the edges in C , listed in clockwise order. Let $a \in V(Y)$ be the vertex that appears in the edge labels $\psi(e_1)$ and $\psi(e_2)$. Because $\{a\}$ is not a dominating set of Y , there must be some $1 \leq i \leq k$ such that $a \notin \psi(e_i)$. (Indeed, if we imagine the vertices of Y sitting on those of Cycle_n , then having $a \in \psi(e_i)$ for all $1 \leq i \leq k$ would mean that the vertex a slides all the way around Cycle_n , participating in friendly swaps with all other vertices of Y .) Let r be the smallest positive integer such that $a \notin \psi(e_r)$. By definition of a , we have $r \geq 3$. Let $\psi(e_{r-2}) = ab$ and $\psi(e_{r-1}) = ac$. Let σ be the vertex of $\text{FS}(\text{Cycle}_n, Y)$ in both the edges e_{r-1} and e_r . Then the two vertices of Cycle_n adjacent to $\sigma^{-1}(c)$ are $\sigma^{-1}(a)$ and $\sigma^{-1}(b)$. This means that we must have $\psi(e_r) = bc$. Thus, the vertices a, b, c form a triangle in Y . Consider the walk in $\text{FS}(\text{Cycle}_n, Y)$ that uses the edges e_{r-2}, e_{r-1}, e_r . We can perform a Yang–Baxter move on this walk to obtain a new walk with three different edges e'_{r-2}, e'_{r-1}, e'_r . The edges $e_{r-2}, e_{r-1}, e_r, e'_r, e'_{r-1}, e'_{r-2}$ form a 6-cycle in $\text{FS}(\text{Cycle}_n, Y)$. ■

Finally, we can prove Theorem 8 by piecing together the results in this section using the same idea as in [9].

Proof of Theorem 8. Suppose Y has domination number at least 3. We will prove the theorem when Y is not triangle-free; the proof when Y is triangle-free is completely analogous (we simply invoke the second statement in Lemma 26 instead of the first statement). Let C be a cycle in $\text{FS}(\text{Cycle}_n, Y)$. We will show that C can be written as a symmetric difference of 4-cycles and 6-cycles. Since friends-and-strangers graphs are bipartite, C has an even number of edges, say $2m$. If $m \in \{2, 3\}$, then we are done. Therefore, we may assume $m \geq 4$ and proceed by induction on m . By induction, we just need to show that C can be written as a symmetric difference of cycles of length strictly smaller than $2m$.

It is straightforward to see that a non-isometric cycle in a graph can be written as a symmetric difference of strictly smaller cycles. Hence, we may assume C is isometric. Let e_1, \dots, e_{2m} be the edges of C , listed in clockwise order. If there are two consecutive edges of C contained in a common 4-cycle, then let $k = 2$; otherwise, let $k = 3$. According to Lemma 26, there are at least k consecutive edges in C that belong to a common $2k$ -cycle; since C is isometric, there must be exactly k edges in C in this $2k$ -cycle. Without loss of generality, say these edges are e_1, \dots, e_k , and let e'_1, \dots, e'_k be the other edges of the common $2k$ -cycle D . Assume that these edges are named so that $e_1, \dots, e_k, e'_k, \dots, e'_1$ is the cyclic order that these edges appear around D . Our choice of k guarantees that D is an isometric cycle of $\text{FS}(\text{Cycle}_n, Y)$. Therefore, Proposition 25 tells us that $\psi(e_i) = \psi(e'_{k+1-i})$ for all $1 \leq i \leq k$. Since C has length at least 8, we can use Proposition 25 again to see that e'_1, \dots, e'_k are not edges in C . Let C' be the cycle with edges $e'_1, \dots, e'_k, e_{k+1}, \dots, e_{2m}$. Then $C = C' \triangle D$, so it suffices to show that C' can be written as a symmetric difference of cycles of length strictly smaller than $2m$. Since C' has $2m$ edges, we just need to prove that C' is not isometric. However, this follows from Proposition 25 because $\psi(e'_k) = \psi(e_1) = \psi(e_{m+1})$ and the edges e'_k and e_{m+1} are not opposite edges of C' . ■

Acknowledgements

We are grateful to the MIT PRIMES organizers for creating such a rare and amazing math research opportunity. Colin Defant was supported by the National Science Foundation under Award No. DGE-1656466 and Award No. 2201907, by a Fannie and John Hertz Foundation Fellowship, and by a Benjamin Peirce Fellowship at Harvard University.

REFERENCES

- [1] N. Alon, C. Defant and N. Kravitz, *Typical and extremal aspects of friends-and-strangers graphs*, J. Combin. Theory Ser. B **158** (2023) 3–42.
<https://doi.org/10.1016/j.jctb.2022.03.001>

- [2] A. Balitskiy and J. Wellman, *Flip cycles in plabic graphs*, Selecta Math. **26** (2020) 15.
<https://doi.org/10.1007/s00029-020-0544-1>
- [3] K. Bangachev, *On the asymmetric generalizations of two extremal questions on friends-and-strangers graphs*, European J. Combin. **104** (2022) 103529.
<https://doi.org/10.1016/j.ejc.2022.103529>
- [4] C. Defant, *Toric promotion*, Proc. Amer. Math. Soc. **151** (2023) 45–57.
<https://doi.org/10.1090/proc/16079>
- [5] C. Defant and N. Kravitz, *Friends and strangers walking on graphs*, Comb. Theory **1** (2021).
<https://doi.org/10.5070/C61055363>
- [6] S. Felsner, L. Kleist, T. Mütze and L. Sering, *Rainbow cycles in flip graphs*, SIAM J. Discrete Math. **34** (2020) 1–39.
<https://doi.org/10.1137/18M1216456>
- [7] R. Jeong, *Diameters of connected components of friends-and-strangers graphs are not polynomially bounded* (2022).
arXiv:2201.00665
- [8] R. Jeong, *On structural aspects of friends-and-strangers graphs* (2022).
arXiv:2203.10337
- [9] M. Naatz, *The graph of linear extensions revisited*, SIAM J. Discrete Math. **13** (2000) 354–369.
<https://doi.org/10.1137/S0895480199352609>
- [10] R.P. Stanley, *An equivalence relation on the symmetric group and multiplicity-free flag h -vectors*, J. Comb. **3** (2012) 277–298.
<https://doi.org/10.4310/JOC.2012.v3.n3.a2>
- [11] L. Wang and Y. Chen, *Connectivity of friends-and-strangers graphs on random pairs*, Discrete Math. **346(3)** (2023) 113266.
<https://doi.org/10.1016/j.disc.2022.113266>
- [12] R.M. Wilson, *Graph puzzles, homotopy, and the alternating group*, J. Combin. Theory Ser. B **16** (1974) 86–96.
[https://doi.org/10.1016/0095-8956\(74\)90098-7](https://doi.org/10.1016/0095-8956(74)90098-7)

Received 15 October 2022

Revised 11 February 2023

Accepted 17 February 2023

Available online 4 April 2023