

CONNECTIVITY OF OLD AND NEW MODELS OF FRIENDS-AND-STRANGERS GRAPHS

ALEKSA MILOJEVIĆ

ABSTRACT. In this paper, we investigate the connectivity of friends-and-strangers graphs, which were introduced by Defant and Kravitz in 2020. We begin by considering friends-and-strangers graphs arising from two random graphs and consider the threshold probability at which such graphs attain maximal connectivity. We slightly improve the lower bounds on the threshold probabilities, thus disproving two conjectures of Alon, Defant and Kravitz. We also improve the upper bound on the threshold probability in the case of random bipartite graphs, and obtain a tight bound up to a factor of $n^{o(1)}$. Further, we introduce a generalization of the notion of friends-and-strangers graphs in which vertices of the starting graphs are allowed to have multiplicities and obtain generalizations of previous results of Wilson and of Defant and Kravitz in this new setting.

1. INTRODUCTION

The main objects of study in this paper are the friends-and-strangers graphs, which were introduced by Defant and Kravitz in 2020 [6].

Definition 1.1. Given two simple graphs X and Y on n vertices, we define the *friends-and-strangers graph* $\text{FS}(X, Y)$ associated to X and Y as follows. Vertices of $\text{FS}(X, Y)$ are all bijections $\sigma : V(X) \rightarrow V(Y)$, and the edges are given by pairs of bijections σ and τ which satisfy $\sigma = \tau \circ (a b)$, for some $a, b \in V(X)$ with $ab \in E(X)$ and $\sigma(a)\sigma(b) \in E(Y)$. In this case, we say that σ and τ differ by a *(X, Y) -friendly swap across the edge $\sigma(a)\sigma(b)$* (or, equivalently, across the edge ab).

The graph $\text{FS}(X, Y)$ has a natural interpretation, from which the name *friends-and-strangers* originated. Namely, if we think of vertices of X as n different people and vertices of Y as n chairs, the bijections $\sigma : V(X) \rightarrow V(Y)$ represent arrangements of people onto chairs, with exactly one person sitting on each chair. We say that a pair of people are friends if they are adjacent in X , and that they are strangers otherwise. Then, a (X, Y) -friendly swap in $\text{FS}(X, Y)$ corresponds to choosing two friends sitting on adjacent chairs and swapping their positions. It is natural to ask whether any arrangement can be obtained from any other by a sequence of friendly swaps, which corresponds to the graph $\text{FS}(X, Y)$ being connected.

The concept of friends-and-strangers graphs generalizes several previously considered problems, as noted in [1] and [6]. For example, in the famous 15-puzzle numbers $1, \dots, 15$ are placed on a 4×4 grid, leaving one cell empty. Any number adjacent to the empty cell is allowed to move to the empty cell, and the goal is to reach a predetermined configuration of numbers. In the context of friends-and-strangers graph, this puzzle corresponds to $\text{FS}(\text{Star}_{16}, \text{Grid}_{4 \times 4})$, the numbers $1, \dots, 15$ corresponding to leaves of Star_{16} and the empty cell corresponding to its center. Generalizing the 15-puzzle, in 1974 Wilson [13] characterized all graphs X for which $\text{FS}(\text{Star}_n, X)$ is connected. Further, Stanley investigated the connected components of $\text{FS}(\text{Path}_n, \text{Path}_n)$ in [12], while Reidys [11] used the graph $\text{FS}(\text{Path}_n, \overline{X})$ in order to investigate the acyclic orientations of X .

Defant and Kravitz [6] recognized that all of these results can be phrased in a unified framework, thus defining the friends-and-strangers graphs. After proving some of the basic properties of these graphs, such as symmetry $\text{FS}(X, Y) \cong \text{FS}(Y, X)$ and the fact that $\text{FS}(X, Y)$ is bipartite, they went on to describe the connected components of $\text{FS}(X, \text{Path}_n)$ and $\text{FS}(X, \text{Cycle}_n)$ in terms of acyclic orientations of the complement \overline{X} . As a consequence of the structural description of $\text{FS}(X, \text{Cycle}_n)$, they showed that $\text{FS}(X, \text{Cycle}_n)$ is connected if and only if the complement of X is a forest of trees of coprime sizes (Corollary 4.14 of [6]). Building on their work, Jeong [8] showed that, for such graphs X , the friends-and-strangers graph $\text{FS}(X, Y)$ is still connected even if Cycle_n is replaced by any other biconnected graph Y . In a follow-up paper, Jeong [9] investigates the diameter of the connected components of the graph $\text{FS}(X, Y)$. Finally, Defant, Dong, Lee, and Wei [5] extended the connectivity results to new classes of graphs, such as spiders and dandelions. Although most of these results address structural questions, other aspects of friends-and-strangers graphs are also worth investigating.

In a follow-up paper Alon, Defant, and Kravitz [1] initiated the study of extremal questions about friends-and-strangers graphs. One of the main questions addressed in [1] asks what is the smallest integer d_n such that for all graphs X, Y whose minimum degrees satisfy $\delta(X), \delta(Y) \geq d_n$, the friends-and-strangers graph $\text{FS}(X, Y)$ is

connected. After Alon, Defant, and Kravitz [1] initially showed $d_n \geq \frac{2}{5}n - 1$, and conjectured $d_n = \frac{2}{5}n + O(1)$, Bangachev [2] settled this conjecture by showing that it is sufficient to have $\delta(X), \delta(Y) > n/2$ and $2 \min\{\delta(X), \delta(Y)\} + 3 \max\{\delta(X), \delta(Y)\} \geq 3n$ for $\text{FS}(X, Y)$ to be connected.

Further, the authors of [1] also introduced several probabilistic questions about the connectivity of $\text{FS}(X, Y)$ and the first half of our paper is dedicated to answering them. More precisely, if X, Y are random graphs, chosen according to the Erdős-Rényi model $\mathcal{G}(n, p)$, Alon, Defant and Kravitz investigated the threshold probability p_{gen} at which $\text{FS}(X, Y)$ becomes connected with high probability. In their paper [1], they showed the inequalities $\Omega(n^{-1/2}) \leq p_{\text{gen}} \leq n^{-1/2+o(1)}$ and conjectured $p_{\text{gen}} = \Theta(n^{-1/2})$. In this paper, we disprove this conjecture by showing that $p_{\text{gen}} \geq \Omega(\log^{1/2} n/n^{1/2})$, using a slightly tweaked version of the original argument from [1]. Asymmetric variant of this question has also been investigated by Wang and Chen [14], who considered the connectivity of $\text{FS}(X, Y)$ when $X \sim \mathcal{G}(n, p_1)$ and $Y \sim \mathcal{G}(n, p_2)$, for different probabilities p_1, p_2 .

Finally, Alon, Defant, and Kravitz studied the case when $X, Y \sim \mathcal{G}(K_{n,n}, p)$ are random bipartite graphs, generated as edge-subgraphs of the balanced bipartite graph $K_{n,n}$ in which every edge is included with probability p , randomly and independently. When X, Y are bipartite graphs, a parity obstruction prevents $\text{FS}(X, Y)$ from being connected and hence it is natural to ask for the threshold probability p_{bip} at which the graph $\text{FS}(X, Y)$ has exactly two components with high probability. In [1], the authors showed that $\Omega(n^{-1/2}) \leq p_{\text{bip}} \leq n^{-0.3+o(1)}$ and conjectured $p_{\text{bip}} = \Theta(n^{-1/2})$, as in the non-bipartite case. We tighten the upper bound by showing that $p_{\text{bip}} \leq n^{-1/2+o(1)}$ using a modified version of the approach from the non-bipartite case, and we tighten the lower bound to $p_{\text{bip}} \geq \Omega(\log^{1/2} n/n^{1/2})$. These results close the polynomial gap between the upper and the lower bound on the threshold probability in the bipartite case, bringing it down to a factor of $n^{o(1)}$.

In Section 2 we present general properties of the $\text{FS}(X, Y)$, some of which reproduced from [1], [6]. Alongside these, in Subsection 2.2 we prove the following strengthening of Propositions 3.1 and 4.1 from [1].

Theorem 1.2. *There exists a constant $\varepsilon > 0$ with the following property. For a large positive integer n and*

$$p < \varepsilon \left(\frac{\log n}{n} \right)^{1/2},$$

if we choose random graphs $X, Y \sim \mathcal{G}(n, p)$ independently, the resulting friends-and-strangers graph $\text{FS}(X, Y)$ has an isolated vertex with high probability. The same statement holds if we choose X, Y to be random bipartite graphs $X, Y \sim \mathcal{G}(K_{n,n}, p)$.

As previously mentioned, Theorem 1.2 disproves Conjectures 7.1 and 7.2 from [1]. Our argument is very similar to the argument given in [1], and the improvement comes from replacing the application of the general Sauer-Spencer theorem by a stronger result of Bollobás, Janson and Scott [4] which applies only in the case of random graphs (which is all we need). Then, in Section 3, we show the following strengthening of Theorem 1.2 from [1].

Theorem 1.3. *Let n be a positive integer, and let X, Y be random bipartite graphs independently chosen from $\mathcal{G}(K_{n,n}, p)$. If*

$$(1) \quad p \geq \frac{\exp(10(\log n)^{4/5})}{n^{1/2}},$$

then $\text{FS}(X, Y)$ has two connected components with high probability.

Remark 1.4. If X and Y are bipartite graphs, Proposition 2.5 of [1] shows that $\text{FS}(X, Y)$ has at least two connected components, due to a parity obstruction. Thus, Theorem 1.3 shows that if p satisfies (1), then the number of connected components of $\text{FS}(X, Y)$ is the least possible. Furthermore, the parity obstruction implies that, if $\text{FS}(X, Y)$ has exactly two connected components, they must have the same size. Thus, if $\text{FS}(X, Y)$ contains an isolated vertex, in the regime of Theorem 1.2, it must have more than two connected components.

Together with Proposition 4.1 of [1], our result determines that $p \sim n^{-1/2}$ is the correct exponent for the threshold probability in the case of $X, Y \sim \mathcal{G}(K_{n,n}, p)$, which agrees with the threshold probability for $\text{FS}(X, Y)$ to be connected when $X, Y \sim \mathcal{G}(n, p)$. Although the proof of Theorem 1.3 follows the general strategy used by Alon, Defant, and Kravitz in the non-bipartite setting, several new ideas are needed, including an additional result about matchings in random bipartite graphs.

In Sections 4 and 5, we investigate a generalization of the friends-and-strangers graphs obtained by adding multiplicities onto vertices. To understand and motivate this generalization, it is useful to adopt an alternative view of friends-and-strangers graphs. We will think of bijections $\sigma : V(X) \rightarrow V(Y)$ as assignments of labels $\sigma(v) \in V(Y)$ to vertices $v \in V(X)$. We will first present a motivating example based on the previous work of Kornhauser, Miller, and Spirakis [10] and then give a formal definition.

Following the early work of Wilson on understanding the graph $\text{FS}(X, \text{Star}_n)$, Kornhauser, Miller, and Spirakis [10] considered the setting in which k different pebbles are placed on the graph with n vertices, and a pebble is allowed to move to any empty vertex adjacent to its current position. In this setting, the result of Wilson [13] can be interpreted as the case $n = k + 1$, when only one vertex is left empty. This scenario can be interpreted in the language of friends-and-strangers graphs by setting X to be the underlying graph on which the pebbles move, and setting the label graph Y to be the graph consisting of k pebbles connected to one blank vertex, in which the blank label is allowed to appear $n - k$ times on X . When $n > k + 1$, the problems of this type do not fit into the framework of friends-and-strangers graphs anymore, since the empty vertices are allowed to repeat. Motivated by this example, let us now formally introduce multiplicity friends-and-strangers graphs.

Definition 1.5. Let X, Y be simple graphs, on n and m vertices, respectively, and let $m \leq n$. Fix a list of multiplicities $c = (c_1, \dots, c_m) \in \mathbb{Z}_{>0}^m$. We call simple graph Y together with the list of multiplicities c a *multiplicity graph*, with *total multiplicity* $\sum_{y \in V(Y)} c_y$. If the total multiplicity of Y is n , we define a multiplicity friends-and-strangers graph $\text{FS}_m(X, Y)$ whose vertices are functions $\sigma : V(X) \rightarrow V(Y)$, called *arrangements*, satisfying $|\sigma^{-1}(\{y\})| = c_y$ for all $y \in V(Y)$. Further, arrangements σ and τ are adjacent in $\text{FS}_m(X, Y)$ when $\sigma = \tau \circ (a b)$, where $ab \in E(X)$ and $\sigma(a)\sigma(b) \in E(Y)$.

Remark 1.6. If the multiplicities of all vertices of Y are equal to 1, we recover the standard friends-and-strangers graph $\text{FS}(X, Y)$. Further, note that the definition of $\text{FS}_m(X, Y)$ is not symmetric in X, Y anymore.

Intuitively, the multiplicity c_y of a label $y \in V(Y)$ can be understood as the number of times this label appears on X . Alternatively, in the interpretation of $\text{FS}(X, Y)$ which involves friends sitting on chairs, the multiplicity c_y can be understood as the number of people from X sitting on a chair $y \in V(Y)$.

Multiplicity friends-and-strangers graphs were also considered in a more applied setting, by Brailovskaya, Gowri, Yu, and Winfree [3]. They proposed a computation model in which various types of molecules are placed on a surface modeled by a grid, and certain pairs of molecules are allowed to swap through a chemical reaction. It turns out that their model can be understood in the language of friends-and-strangers graphs, the molecules types corresponding to labels placed on the grid graph, with the edges between molecule types encoding which swaps are chemically allowed.

In Section 4, we focus on investigating connectivity of the graphs $\text{FS}_m(X, \text{Star}_n)$ and $\text{FS}_m(\text{Star}_n, X)$, generalizing previous work of Wilson. Before formally stating our results, let us present the precise statement of Wilson’s theorem.

Theorem 1.7 ([13]). *Let X be a simple graph on $n \geq 3$ vertices. The graph $\text{FS}(X, \text{Star}_n)$ is connected if and only if X is a biconnected graph, which is not bipartite, not a cycle of length at least 4 and not isomorphic to θ_0 , which a specific graph on 7 vertices shown on Figure 1. Furthermore, if X is bipartite and biconnected, but it is not a cycle of length at least 4 or θ_0 , then the graph $\text{FS}(X, \text{Star}_n)$ has exactly two connected components.*

If a graph X satisfies all conditions of Wilson’s theorem for $\text{FS}(X, \text{Star}_n)$ to be connected, we say X is *Wilsonian*.




FIGURE 1. Diagram of the graph θ_0 .

Wilson’s theorem can be generalized in two possible ways, depending if the multiplicities are assigned to the graph X or to the star Star_n . Theorem 1.8 characterizes the sets of multiplicities that can be assigned to vertices of X such that $\text{FS}_m(\text{Star}_n, X)$ is still connected, while Theorem 1.9 does the same in the case the multiplicities are assigned to the vertices of Star_n . Theorem 1.9 was originally derived by Kornhauser, Miller, and Spirakis in 1984 [10], using algebraic methods. Since some of the details of their proof are omitted, we present a rigorous, purely combinatorial proof of their results in the language of friends-and-strangers graphs.

Theorem 1.8. *Let X be a connected multiplicity graph on the vertex set $[m]$ with multiplicities $c = (c_1, \dots, c_m) \in \mathbb{Z}_{>0}^m$. If X is Wilsonian, the graph $\text{FS}_m(\text{Star}_n, X)$ is connected. If X is biconnected but not Wilsonian, the graph $\text{FS}_m(\text{Star}_n, X)$ is connected if and only if there exists a vertex $v \in V(X)$ with multiplicity $c_v \geq 2$. Finally, if X is not biconnected, the graph $\text{FS}_m(\text{Star}_n, X)$ is connected if and only if $c_v \geq 2$ for all cut vertices $v \in V(X)$.*

To state Theorem 1.9, we need to introduce the following notion. Given a simple graph X , a k -bridge in X is a k -tuple of vertices $a_1, \dots, a_k \in V(X)$ for which the following two conditions hold: a_{i+1} and a_{i-1} are the only neighbors of a_i , for $i \in \{2, \dots, k-1\}$, and in $X|_{V(X) - \{a_2, \dots, a_{k-1}\}}$ the vertices a_1 and a_k lie in different connected components, which both have size at least 2.



FIGURE 2. Diagram of a k -bridge consisting of vertices a_1, \dots, a_k .

Theorem 1.9 ([10]). *Let Star_m be a star on $m > 2$ vertices and let $c \in \mathbb{Z}_{>0}^m$ be the list of multiplicities assigned to vertices of Star_m , with the center of Star_m having multiplicity $k \geq 2$. Let X be a connected simple graph on n vertices, where n is the total multiplicity of Star_m . If X does not contain a k -bridge and if X is not a cycle, the graph $\text{FS}_m(X, \text{Star}_m)$ is connected. If X contains a k -bridge, $\text{FS}_m(X, \text{Star}_m)$ is disconnected. If X is a cycle, the graph $\text{FS}_m(X, \text{Star}_m)$ is connected if and only if $m = 3$ and one of the leaves of Star_m has multiplicity 1.*

In Section 5, we describe the structure of $\text{FS}_m(\text{Path}_n, X)$ and $\text{FS}_m(\text{Cycle}_n, X)$ when X is a multiplicity graph. In [6], Defant and Kravitz characterized the connected components of the graphs $\text{FS}(\text{Path}_n, Y)$, $\text{FS}(\text{Cycle}_n, Y)$ in terms of the acyclic orientations of the complement \overline{Y} , using the notion of *toric equivalence*, which was initially introduced by Develin, Macauley, and Reiner in [7]. Following this approach, we describe the connected components of $\text{FS}_m(\text{Path}_n, X)$, $\text{FS}_m(\text{Cycle}_n, X)$ in terms of the acyclic orientations of $\overline{X'}$, where X' denotes the blow-up of X with respect to the multiplicity list c . We then use this description to characterize all graphs X and multiplicity lists c for which $\text{FS}_m(\text{Cycle}_n, X)$ is connected. Since formal definitions and exact statements of these results are somewhat technical, we postpone them until Section 5.

In Section 6, we end by presenting another generalization of friends-and-strangers graphs, which allows both graphs X and Y to have multiplicities, along with indicating possible further research directions.

2. PRELIMINARIES

In this section, we present some of the basic tools we will use throughout the paper. We begin by presenting two results used to show friends-and-strangers graphs are connected, and then show two typical obstructions to connectivity of these graphs. Together with these obstructions, we discuss the proof of Theorem 1.2 which shows that certain random friends-and-strangers graphs have isolated vertices. Finally, we will conclude by introducing the *lift* of a multiplicity graph, a notion that proves to be key in understanding this generalization of friends-and-strangers graphs.

2.1. Connectivity of $\text{FS}(X, Y)$. One of the fundamental notions Alon, Defant, and Kravitz use in [1] to show their results about connectivity of friends-and-strangers graphs are exchangeable pairs. More precisely, for a bijection $\sigma : V(X) \rightarrow V(Y)$, we say that a pair of vertices $u, v \in V(Y)$ is (X, Y) -exchangeable from σ if there exists a sequence of (X, Y) -friendly swaps which transforms the arrangement σ into $(u v) \circ \sigma$. Note that this sequence might consist of more than one swap. Since $\text{FS}(X, Y)$ is symmetric in X and Y , one might define the notion of exchangeability for pairs of vertices $u', v' \in V(X)$ in the symmetric fashion. The following result, stated as Proposition 2.8 in [1], shows how exchangeability can be used to establish connectivity of $\text{FS}(X, Y)$.

Proposition 2.1 ([1]). *Let X, Y, \tilde{Y} be n -vertex graphs and suppose $Y \subseteq \tilde{Y}$. Suppose that for every edge uv of \tilde{Y} and every bijection $\sigma : V(X) \rightarrow V(Y)$ satisfying $\sigma^{-1}(u)\sigma^{-1}(v) \in E(X)$, the vertices u and v are (X, Y) -exchangeable from σ . Then, the connected components of $\text{FS}(X, Y)$ and the connected components of $\text{FS}(X, \tilde{Y})$ have the same vertex sets. In particular, the number of connected components of $\text{FS}(X, \tilde{Y})$ is equal to the number of connected components of $\text{FS}(X, Y)$.*

Note that Proposition 2.1 allows us to translate a global question about connectedness of $\text{FS}(X, Y)$ to a local question about exchangeability of certain pairs of vertices, and therefore it will be very useful throughout the paper. Since the graph $\text{FS}(X, K_n)$ is connected whenever X is connected, the following special case of Proposition 2.1, stated as Lemma 2.9 in [1], is particularly interesting. Setting $\tilde{Y} = K_n$ shows that, for a connected graph X , the graph $\text{FS}(X, Y)$ is connected if every pair of adjacent vertices in X is (X, Y) -exchangeable from any arrangement σ .

In the setting of multiplicity graphs, the notion (X, Y) -exchangeability generalizes with almost the same definition. The only difference is that exchangeability can only be defined for pairs of vertices of X . Formally, we

say that a pair of vertices $u', v' \in V(X)$ is (X, Y) -exchangeable from a bijection σ if there exists a sequence of (X, Y) -friendly swaps transforming σ into $\sigma \circ (u' v')$. It is not hard to see that the analogue of Proposition 2.1 holds even in the setting of multiplicity graphs.

The following useful result, first observed by Bangachev in [2], relates the notion of exchangeability from two different arrangements, σ and σ' .

Proposition 2.2 ([2]). *Suppose that $\sigma, \sigma' \in V(\text{FS}_m(X, Y))$ are in the same connected component of $\text{FS}_m(X, Y)$ and assume that the sequence of swaps transforming σ into σ' never involves vertices $u, v \in V(X)$. Then, if the pair of vertices u, v is (X, Y) -exchangeable from σ , it is also (X, Y) -exchangeable from σ' .*

Proof. Suppose that Σ is the sequence of (X, Y) -friendly swaps which transform σ into σ' , and let Σ^{-1} be the reverse sequence which transforms σ' into σ . Since none of the swaps in Σ involve u, v , it is simple to see that Σ also transforms $\sigma \circ (u v)$ into $\sigma' \circ (u v)$. Given that the pairs (σ', σ) , $(\sigma, \sigma \circ (u v))$ and $(\sigma \circ (u v), \sigma' \circ (u v))$ are in the same connected component of $\text{FS}_m(X, Y)$, we conclude that the pair u, v is (X, Y) -exchangeable from σ' . \square

2.2. Obstructions to connectivity of friends-and-strangers graphs. We will now present two typical obstructions to connectivity of $\text{FS}(X, Y)$. The first one occurs when both X and Y have cut vertices, where a cut vertex is defined as vertex whose removal disconnects the graph. The obstruction presented in the following proposition is a direct generalization of Proposition 2.6 from [6] to the setting with multiplicities. Since the proof is completely analogous to the original, we omit it.

Proposition 2.3. *Let X be a connected graph on $m \geq 3$ vertices with a list of multiplicities $c \in \mathbb{Z}_{>0}^m$, which has total multiplicity n , and let Y be a graph on n vertices. Suppose $x_0 \in V(X)$, $y_0 \in V(Y)$ are cut vertices and that $c_{x_0} = 1$. Also, let X_1, \dots, X_r be the connected components of X produced by removing x_0 , and let Y_1, \dots, Y_s be the connected components of Y produced by removing y_0 . Let M be the set of $r \times s$ matrices with nonnegative integer entries in which the i -th row sums to $|X_i|$ and the j -th column sums to $|Y_j|$. Then, $\text{FS}_m(Y, X)$ has at least $|M|$ connected components.*

In case Y has at least 3 vertices, it is not hard to see that $|M| \geq 2$, and therefore Proposition 2.3 implies that the graph $\text{FS}_m(Y, X)$ is disconnected.

Let us now present the second obstruction to connectivity which will be important in this paper. It was mentioned in the introduction that if X, Y are bipartite graphs, the graph $\text{FS}(X, Y)$ has at least two connected components, due to a parity obstruction. Moreover, if the graph $\text{FS}(X, Y)$ has exactly two connected components, there is a simple way to test whether two arrangements are in the same connected component, which is described by the following proposition. This description will be particularly useful in Section 3, where we consider random bipartite graphs X and Y . This proposition was presented in [1] as Proposition 2.5 and it is equivalent to Proposition 2.7 of [6].

Proposition 2.4 ([1], [6]). *Let X and Y be bipartite graphs on n vertices with vertex bipartitions $V(X) = A_X \sqcup B_X$ and $V(Y) = A_Y \sqcup B_Y$. If the bijections σ and τ are in the same connected component of $\text{FS}(X, Y)$, then $\text{sgn}(\sigma^{-1} \circ \tau)$ has the same parity as $|\tau(A_X) \cap A_Y| - |\sigma(A_X) \cap A_Y|$.*

Here, the sign of the permutation π , denoted by $\text{sgn}(\pi)$ denotes the parity of the number of inversions in the permutation π . Finally, we present the proof of Theorem 1.2 closely the approach from [1].

Proof of Theorem 1.2. We begin by focusing on the non-bipartite case. Given two n -vertex graphs X, Y , an isolated vertex in the graph $\text{FS}(X, Y)$ corresponds to a bijection $\sigma : V(X) \rightarrow V(Y)$ with the property that $\sigma(u)\sigma(v) \notin E(Y)$ for all edges $uv \in E(X)$. In other words, an isolated vertex of $\text{FS}(X, Y)$ corresponds to a packing of graphs X, Y . Theorem 1 from the paper of Bollobás, Janson and Scott [4] gives a criterion for the existence of such packing, when X and Y are random graphs generated from $G(n, p)$. More precisely, plugging in $k = 2$ and $p = q$ into this theorem gives the following statement: there exists a constant $\varepsilon > 0$ such that for $p^2 \leq \varepsilon \log n/n$ and independent random graphs $X, Y \sim G(n, p)$, there is a packing of X, Y in K_n with high probability. In other words, for $p < (\varepsilon \log n/n)^{1/2}$ the friends-and-strangers graph $\text{FS}(X, Y)$ contains an isolated vertex.

The case of bipartite graphs is almost identical. Namely, random bipartite graphs on $2n$ vertices can be generated by taking $G(2n, p)$ and deleting all edges occurring among the first n vertices and among the last n vertices. As long as $p < (\varepsilon \log 2n/2n)^{1/2}$, two independent graphs from $G(2n, p)$ can be packed with high probability, and hence the same holds after deleting the edges. This completes the proof of Theorem 1.2. \square

Remark 2.5. Since the bounds of Theorem 1 in [4] are tight, this proof actually determines exactly the threshold probability for the existence of isolated vertices in randomly generated friends-and-strangers graphs.

2.3. Lift of a multiplicity graph. Finally, the key tool for understanding the connectivity of multiplicity friends-and-strangers graphs is the lift of a multiplicity graph.

Definition 2.6. Let Y be a graph on m vertices and let $c = (c_1, \dots, c_m) \in \mathbb{Z}_{>0}^m$ be its multiplicity list. We define the *lift* of Y to be a simple graph Y' on $\sum_{v \in V(Y)} c_v$ vertices which has a vertex partition of Y' , $V(Y') = \bigsqcup_{v \in V(Y)} S_v$, with the following properties:

- $Y'|_{S_v}$ is a clique of size c_v for all vertices $v \in V(Y)$, and
- vertices $u' \in S_u, v' \in S_v$ are adjacent in Y' if and only if u, v are adjacent in Y .

In other words, the lift Y' is a blow-up of the graph Y , in which the vertex v is replaced with a clique of size c_v . Note that defining the lift of the graph also defines a natural projection map π_Y from Y' to Y which maps all elements of S_v to v .

To see why considering the lift of a graph is useful, let us describe an equivalence relation on the vertices of $\text{FS}(X, Y')$, along with a way to describe the structure of $\text{FS}_m(X, Y)$ using $\text{FS}(X, Y')$. Throughout this discussion, we assume a multiplicity graph Y is given, with total capacity n , together with a simple graph X on n vertices.

For a finite set A , let \mathfrak{S}_A denote the group of all permutations $\rho : A \rightarrow A$, with the composition operation. Furthermore, for a multiplicity graph Y , let us define \mathfrak{S}_Y to be the group of permutations $\rho : V(Y') \rightarrow V(Y')$ which satisfy $\rho(S_v) = S_v$ for all cliques S_v . In other words, \mathfrak{S}_Y is the set of those permutations that permute the vertices of Y' only within their cliques. Then, we define an equivalence relation \equiv on the bijections from $V(X)$ to $V(Y')$ given by $\sigma_1 \equiv \sigma_2$ if $\sigma_1 = \rho \circ \sigma_2$ for some $\rho \in \mathfrak{S}_Y$. Equivalently, one might say that $\sigma_1 \equiv \sigma_2$ if $\pi_Y \circ \sigma_1 = \pi_Y \circ \sigma_2$, as functions from $V(X)$ to $V(Y)$.

This equivalence relation induces a natural quotient operation on the graph $\text{FS}(X, Y')$, which identifies all vertices equivalent under \equiv and assigns an edge between distinct equivalence classes $[\sigma]_{\equiv}, [\tau]_{\equiv}$ precisely when there exist representatives $\sigma_0 \in [\sigma]_{\equiv}, \tau_0 \in [\tau]_{\equiv}$ which are adjacent in $\text{FS}(X, Y')$. We denote the arising graph by $\text{FS}(X, Y')/\equiv$.

Proposition 2.7. *Let X be a simple graph on n vertices Y be a multiplicity graph with total multiplicity n . Then, the quotient $\text{FS}(X, Y')/\equiv$ is isomorphic to the graph $\text{FS}_m(X, Y)$.*

Before passing to the formal proof, let us illustrate this statement with an example.

Example 2.8. Suppose that $X = \text{Path}_3$, with vertices numbered 1, 2, 3, and Y is a multiplicity graph on two adjacent vertices, with multiplicities 1 and 2. The lift Y' is a triangle and the graphs $\text{FS}_m(X, Y)$ and $\text{FS}(X, Y')$ are represented schematically in Figure 3.




FIGURE 3. Schematic representation of graphs $\text{FS}_m(X, Y)$ on the left and $\text{FS}(X, Y')$ on the right.

The red contours represent vertices of the graphs $\text{FS}_m(X, Y)$ and $\text{FS}(X, Y')$. Within each red vertex, we see the arrangement of labels 1, 2, 3 from $V(X)$ onto the vertices of Y or Y' corresponding to it. Note that the vertices of Y' are colored in green and blue in order to keep track which vertex of Y they are coming from.

In this example, \mathfrak{S}_Y consists of only two permutations, and the induced equivalence relation \equiv has equivalence classes of size 2. These equivalence classes are precisely the pairs of vertices of $\text{FS}(X, Y')$ which are horizontally next to each other — they differ precisely by exchanging the labels of the two green vertices.

Proof of Proposition 2.7. Let us note that every equivalence class of bijections $[\sigma]_{\equiv}$ can be associated with a function $f_{[\sigma]_{\equiv}} = \pi_Y \circ \sigma$ from $V(X)$ to $V(Y)$, which is independent of the representative $\sigma \in [\sigma]_{\equiv}$. The isomorphism of

$\text{FS}(X, Y')/\equiv$ and $\text{FS}_m(X, Y)$ is given by the correspondence $[\sigma]_{\equiv} \mapsto f_{[\sigma]_{\equiv}}$. In the rest of the proof, we check that this correspondence preserves the edges of $\text{FS}(X, Y')/\equiv$.

We begin by showing that if $[\sigma]_{\equiv}$ and $[\tau]_{\equiv}$ are distinct adjacent vertices of $\text{FS}(X, Y')/\equiv$, then $f_{[\sigma]_{\equiv}}$ and $f_{[\tau]_{\equiv}}$ are also adjacent in $\text{FS}_m(X, Y)$. By definition of the quotient graph $\text{FS}(X, Y')$, we have representatives $\sigma_0 \in [\sigma]_{\equiv}, \tau_0 \in [\tau]_{\equiv}$ differing by a (X, Y') -friendly swap along $u_X v_X \in E(X)$, i.e. $\sigma_0 = \tau_0 \circ (u_X v_X)$. Since σ_0 and τ_0 are in distinct equivalence classes, the vertices $\sigma_0(u_X), \tau_0(v_X) \in V(Y')$ are adjacent, but not in the same clique in the clique decomposition of Y' . Therefore, the vertices $\pi_Y \circ \sigma_0(u_X), \pi_Y \circ \tau_0(v_X)$ are distinct and adjacent in Y , meaning that swapping vertices u_X, v_X is a (X, Y) -friendly swap for the function $f_{[\sigma]_{\equiv}}$. Since $f_{[\sigma]_{\equiv}} = f_{[\tau]_{\equiv}} \circ (u_X v_X)$, the vertices $f_{[\sigma]_{\equiv}}$ and $f_{[\tau]_{\equiv}}$ are adjacent in $\text{FS}(X, Y)$.

To show the converse statement, suppose $f_{[\sigma]_{\equiv}}$ and $f_{[\tau]_{\equiv}}$ are vertices of $\text{FS}(X, Y)$ which differ by a (X, Y) -friendly swap of adjacent labels $u_X, v_X \in V(X)$ along the edge $u_Y v_Y \in E(Y)$, where u_Y, v_Y are vertices of Y . For any representative $\sigma_0 \in [\sigma]_{\equiv}$, we have $\sigma_0(u_X) \in S_{u_Y}$ and $\sigma_0(v_X) \in S_{v_Y}$, which means that $\sigma_0(u_X)\sigma_0(v_X)$ is an edge of Y' . Therefore, the swap $\sigma_0 \mapsto \sigma_0 \circ (u_X v_X)$ is a (X, Y') -friendly swap. Noting that $\sigma_0 \circ (u_X v_X) \equiv \tau$, we conclude that there exists a (X, Y') -friendly swap mapping between representatives of $[\sigma]_{\equiv}$ and $[\tau]_{\equiv}$. This completes the proof of Proposition 2.7. \square

Since the main focus of this work is connectivity of $\text{FS}_m(X, Y)$, the following simple corollary gives us an efficient criterion for determining the connectivity of $\text{FS}_m(X, Y)$ in terms of connectivity of $\text{FS}(X, Y')$.

Corollary 2.9. *Let X be a simple graph on n vertices and let Y be a multiplicity graph of total multiplicity n . If the graph $\text{FS}(X, Y')$ is connected, then so is $\text{FS}_m(X, Y)$.*

Proof. It is immediate to see that the quotient $\text{FS}(X, Y')/\equiv$ is connected whenever $\text{FS}(X, Y')$ is connected. Together with Proposition 2.7, this implies connectedness of $\text{FS}_m(X, Y)$. \square

In fact, we can establish an even more precise description of the connected components of $\text{FS}_m(X, Y)$ just based on $\text{FS}(X, Y')$, which will be very important in Section 5.

Corollary 2.10. *Let X be a simple graph on n vertices and let Y be a multiplicity graph of total multiplicity n . Let $E_0 = \{\sigma\tau \mid \sigma, \tau \in V(\text{FS}(X, Y')) \text{ and } \sigma \equiv \tau\}$ and let G be the graph obtained by adding the edges of E_0 to $\text{FS}(X, Y')$. Then, the connected components of $\text{FS}_m(X, Y)$ have the same vertex sets as projections of connected components of G under the map $\sigma \mapsto \pi_Y \circ \sigma$.*

Proof. By Proposition 2.7, we may consider the graph $\text{FS}(X, Y')/\equiv$ instead of $\text{FS}_m(X, Y)$, with the map $\sigma \mapsto [\sigma]_{\equiv}$ replacing $\sigma \mapsto \pi_Y \circ \sigma$. Let us denote the projection $\sigma \mapsto [\sigma]_{\equiv}$ by π_{\equiv} .

Note that the partition of $V(G)$ into connected components is coarser than the partition into equivalence classes $[\sigma]_{\equiv}$, and therefore different connected components have disjoint images under π_{\equiv} . Since the map $\pi_{\equiv} : V(\text{FS}(X, Y')) \rightarrow V(\text{FS}(X, Y')/\equiv)$ is surjective, to complete the proof it suffices to show that, for an arbitrary connected component C of the graph G , $(\text{FS}(X, Y')/\equiv)|_{\pi_{\equiv}(C)}$ is a connected component of $\text{FS}(X, Y')/\equiv$.

We begin by showing there are no outgoing edges from $\pi_{\equiv}(C)$. Suppose the contrary, and let $[\sigma]_{\equiv}[\tau]_{\equiv}$ be an edge of $\text{FS}(X, Y')/\equiv$ with $[\sigma]_{\equiv} \in \pi_{\equiv}(C), [\tau]_{\equiv} \notin \pi_{\equiv}(C)$. By definition of $\text{FS}(X, Y')/\equiv$, this means there exist representatives $\sigma_0 \in [\sigma]_{\equiv}, \tau_0 \in [\tau]_{\equiv}$ with $\sigma_0\tau_0 \in E(\text{FS}(X, Y'))$. Since the partition of $V(G)$ into connected components is coarser than the partition into equivalence classes, we must have $\sigma_0 \in C$ and $\tau_0 \notin C$, which is a contradiction to the fact that C is a connected component of G .

On the other hand, by definition of $\text{FS}(X, Y')/\equiv$, the map π_{\equiv} preserves adjacencies between nonequivalent bijections, which means that $(\text{FS}(X, Y')/\equiv)|_{\pi_{\equiv}(C)}$ is a connected subgraph of $\text{FS}(X, Y')/\equiv$. This completes the proof of Corollary 2.10. \square

3. RANDOM BIPARTITE GRAPHS

3.1. General proof strategy. The main goal of this section is to prove Theorem 1.3, following the proof strategy presented in [1]. We begin by introducing the terminology used in the proof and presenting a high-level overview of the proof strategy.

The first step of the proof is to reduce the question about connectivity of the graph $\text{FS}(X, Y)$ to a local problem about exchangeability of certain pairs of vertices. More precisely, we begin by reducing Theorem 1.3 to Proposition 3.1, which mirrors Proposition 4.5 of [1].

Proposition 3.1. *Let X, Y be independently chosen random graphs in $\mathcal{G}(K_{n,n}, p)$, where p satisfies (1), as in Theorem 1.3. Then, the following statement holds with high probability: for any arrangement $\sigma : V(X) \rightarrow V(Y)$ and any two vertices u_0, v_0 in different partite sets of Y whose preimages $\sigma^{-1}(u_0), \sigma^{-1}(v_0)$ are in different partite*

sets of X , one has that u_0 and v_0 are (\tilde{X}, Y) -exchangeable from σ , where \tilde{X} denotes the graph X with edge $\sigma^{-1}(u_0)\sigma^{-1}(v_0)$ added.

The following argument, reproduced from [1], shows how Proposition 3.1 implies Theorem 1.3. Intuitively, this argument should be thought of as the analogue of Proposition 2.1 in the bipartite setting.

Proof of Theorem 1.3 from Proposition 3.1. We use Proposition 2.8 from [1], which states the following. Let X, Y and \tilde{Y} be graphs on $[n]$, having $Y \subseteq \tilde{Y}$. Suppose that for every edge $uv \in E(\tilde{Y})$ and every $\sigma : V(X) \rightarrow V(Y)$ with $\sigma^{-1}(u)\sigma^{-1}(v) \in E(X)$ the vertices u, v are (X, Y) -exchangeable from σ . Then the number of connected components of $\text{FS}(X, Y)$ is the same as the number of connected components of $\text{FS}(X, \tilde{Y})$.

Let us begin by applying Proposition 3.1 to conclude that, with high probability, any two vertices u_0, v_0 in different partite sets of Y satisfying $\sigma^{-1}(u_0)\sigma^{-1}(v_0) \in E(X)$ are (X, Y) -exchangeable, as $\tilde{X} = X$. Hence, by setting $\tilde{Y} = K_{n,n}$ in Proposition 2.8 of [1], we conclude that $\text{FS}(X, Y)$ has the same number of connected components as $\text{FS}(X, K_{n,n})$. Since $\text{FS}(X, K_{n,n}) \cong \text{FS}(K_{n,n}, X)$, it suffices to show that $\text{FS}(K_{n,n}, X)$ has two connected components with high probability.

Swapping the roles of X and Y in Proposition 3.1, we can also say that, with high probability, for any two vertices u'_0, v'_0 in different partite sets of X , u'_0 and v'_0 are (\tilde{Y}, X) -exchangeable from any $\tau : V(Y) \rightarrow V(X)$, where \tilde{Y} is the graph Y with $\tau^{-1}(u'_0)\tau^{-1}(v'_0)$ added. Since $\tilde{Y} \subset K_{n,n}$, any two vertices u'_0, v'_0 in different partite sets of X are $(K_{n,n}, X)$ -exchangeable.

Therefore, by Proposition 2.8 of [1], we conclude that $\text{FS}(K_{n,n}, X)$ has the same number of connected components as $\text{FS}(K_{n,n}, K_{n,n})$, which is two, as shown by Proposition 2.6 of [1]. Hence, we conclude that the graph we started with, $\text{FS}(X, Y)$, also has two connected components with high probability, completing the proof. \square

To prove Proposition 3.1, we combine the methods used in [1] to tackle the bipartite and non-bipartite cases. The approach for proving Proposition 3.1 relies on constructing sparse bipartite graphs G, H on the vertex set $[m] \cup \{u, v\}$, in which the vertices u and v are (G, H) -exchangeable starting from the identity arrangement. Since the construction is involved and contains many details, we postpone the details until Subsections 3.2 and 3.3. Still, we will point out the properties of G, H relevant for the proof.

After constructing G and H , our next objective will be to show the following statement: With high probability, for every arrangement σ and vertices u_0, v_0 from Proposition 3.1, there exist embeddings $\psi_G : V(G) \rightarrow V(\tilde{X})$ and $\psi_H : V(H) \rightarrow V(Y)$ respecting the arrangement σ in the sense that $\sigma \circ \psi_G = \psi_H \circ \text{Id}$ and having $\psi_G(u) = u_0, \psi_G(v) = v_0$. In this context, an embedding is merely an injective map which preserves adjacencies. This objective will be accomplished in Subsections 3.4 and 3.5.

Intuitively, finding embeddings ψ_G, ψ_H corresponds to finding subgraphs of X and Y which are isomorphic to G and H and are mapped to each other under σ . Proposition 3.2 shows how embeddings ψ_G, ψ_H can be used to ensure that u_0, v_0 are (X, Y) -exchangeable.

Proposition 3.2. *Let X, Y be bipartite graphs with the vertex set $[n]$, and let $\sigma : V(X) \rightarrow V(Y)$ be an arbitrary arrangement. Further, let G, H be bipartite graphs with the vertex set $[m] \cup \{u, v\}$, such that u and v are (G, H) -exchangeable starting from the identity arrangement $\text{Id} : V(G) \rightarrow V(H)$. Suppose that there exist embeddings $\psi_G : V(G) \rightarrow V(\tilde{X})$ and $\psi_H : V(H) \rightarrow V(Y)$ with $\sigma \circ \psi_G = \psi_H \circ \text{Id}$ and $\psi_G(u) = u_0, \psi_G(v) = v_0$. Then, u_0 and v_0 are (\tilde{X}, Y) -exchangeable from σ .*

Proof. Let the sequence of (G, H) -friendly swaps transforming Id into $(u v)$ be $a_1b_1, a_2b_2, \dots, a_kb_k$. Then, the sequence of swaps $\psi_H(a_1)\psi_H(b_1), \dots, \psi_H(a_k)\psi_H(b_k)$ takes σ to $(u_0 v_0) \circ \sigma$. Properties of ψ_G, ψ_H guarantee that all of these swaps are (\tilde{X}, Y) -friendly, and hence (u_0, v_0) are (\tilde{X}, Y) -exchangeable. \square

Finally, let us describe the strategy we use to construct the embeddings $\psi_G : V(G) \rightarrow V(X)$ and $\psi_H : V(H) \rightarrow V(Y)$. The idea behind this step of the proof is to choose large disjoint sets $V_1, \dots, V_m \subseteq V(Y)$ and look for a m -tuple of vertices $(\psi_H(1), \dots, \psi_H(m)) \in V_1 \times \dots \times V_m$ which will, together with vertices u_0 and v_0 , span a copy of H in Y and whose preimages under σ span a copy of G in X . To formalize this idea, we need to introduce the following terminology, which initially appeared in [1].

Let $G' = G|_{[m]}$ and $H' = H|_{[m]}$ be graphs G, H with vertices u, v removed, and let $V(G') = A_{G'} \sqcup B_{G'}$, $V(H') = A_{H'} \sqcup B_{H'}$ be the respective bipartitions. Also, assume that X and Y have bipartitions $V(X) = A_X \sqcup B_X$ and $V(Y) = A_Y \sqcup B_Y$. For a fixed arrangement $\sigma : V(X) \rightarrow V(Y)$, we call disjoint sets $V_1, \dots, V_m \subseteq V(Y)$ *admissible* for σ and (G', H') if the following two conditions are satisfied:

- $V_i \subseteq A_Y$ for $i \in A_{H'}$ and $V_i \subseteq B_Y$ if $i \in B_{H'}$, and
- $\sigma^{-1}(V_i) \subseteq A_X$ for $i \in A_{G'}$ and $\sigma^{-1}(V_i) \subseteq B_X$ for $i \in B_{G'}$.

We say that the pair (G', H') is *bipartite-embeddable in (X, Y)* with respect to admissible sets V_1, \dots, V_m and an arrangement σ if there exist embeddings $\psi_G : V(G') \rightarrow V(X), \psi_H : V(H') \rightarrow V(Y)$ such that $\psi_H(i) \in V_i$ and $\sigma \circ \psi_G = \psi_H \circ \text{Id}$. Moreover, if q_1, \dots, q_m are fixed positive integers satisfying $q_1 + \dots + q_m \leq n$, we say that (G', H') is *(q_1, \dots, q_m) -bipartite-embeddable in (X, Y)* if (G', H') is bipartite-embeddable in (X, Y) for any admissible choice of V_1, \dots, V_m and σ satisfying $|V_i| = q_i$ for all $i \in [m]$.

The concept of (q_1, \dots, q_m) -bipartite-embeddability is useful for the following reasons. On one hand, Lemma 4.2 of [1] gives an simple condition on (q_1, \dots, q_m) which ensures that the pair (G', H') is (q_1, \dots, q_m) -bipartite-embeddable in X, Y with high probability. Since each of q_1, \dots, q_m must be somewhat large in order to satisfy this condition, we cannot directly infer the existence of embeddings ψ_G, ψ_H , since Proposition 3.2 requires that $\psi_H(u) = u_0$ and $\psi_H(v) = v_0$.

Hence, we need to use the fact (G', H') is (q_1, \dots, q_m) -bipartite-embeddable in the following way. The idea is to choose $V_i \subseteq N_Y(u_0)$ for all $i \in N_H(u)$ and $V_i \subseteq \sigma(N_X(u'_0))$ for all $i \in N_G(u)$, since this ensures $\psi_H(i)u_0 \in E(Y), \psi_G(i)u'_0 \in E(X)$ (we require similar conditions for $i \in N_G(v), N_H(v)$). However, since the sets V_1, \dots, V_m must be admissible, we must simultaneously ensure $V_i \subseteq A_Y$ for all $i \in A_H$ and $V_i \subseteq \sigma(A_X)$ for $i \in A_G$. Hence, based on how the pairs of sets $\sigma(N_X(u'_0)), A_Y$ and $N_Y(u_0), \sigma(A_X)$ intersect, we have four cases. To describe them, we need the notion of *majority-mapping*, introduced in [1].

For an arrangement $\sigma : V(X) \rightarrow V(Y)$ and vertices u_0, v_0 having preimages $u'_0 = \sigma^{-1}(u_0)$ and $v'_0 = \sigma^{-1}(v_0)$, we say that σ *majority-maps* $N_X(u'_0)$ to A_Y if $|\sigma(N_X(u'_0)) \cap A_Y| \geq \frac{1}{2}|N_X(u'_0)|$. This definition is naturally generalized if we replace A_Y by B_Y or if we replace σ and Y by σ^{-1} and X . We adopt the convention of naming A_X and A_Y so that $u'_0 \in A_X$ and $u_0 \in A_Y$. Based on where σ and σ^{-1} majority-maps $N_X(u'_0), N_X(v'_0), N_Y(u_0), N_Y(v_0)$, we have four essentially different cases.¹

- (C1): σ majority-maps $N(u'_0)$ and $N(v'_0)$ to B_Y , and σ^{-1} majority-maps $N(u_0)$ and $N(v_0)$ to A_X .
- (C2): σ majority-maps $N(u'_0)$ and $N(v'_0)$ to B_Y , and σ^{-1} majority-maps $N(u_0)$ to A_X and $N(v_0)$ to B_X .
- (C3): σ majority-maps $N(u'_0)$ to B_Y and $N(v'_0)$ to A_Y , and σ^{-1} majority-maps $N(u_0)$ to A_X and $N(v_0)$ to B_X .
- (C4): σ majority-maps $N(u'_0)$ to B_Y and $N(v'_0)$ to A_Y , and σ^{-1} majority-maps $N(u_0)$ to B_X and $N(v_0)$ to A_X .

As we will see, all other cases can be reduced to those four by applying symmetry. These cases seem to be merely a technical requirement of our argument and it might be possible to apply some additional symmetries, or slightly alter the argument, in order to further reduce the number of cases.

In anticipation of these four cases, we will construct four pairs of graphs $(G^{(\rho)}, H^{(\rho)})$, for $\rho \in \{1, 2, 3, 4\}$, in which sets $N_G(u), N_H(u), N_G(v), N_H(v)$ are disjoint and satisfy the following properties, corresponding to the cases (C1), (C2), (C3), (C4).

- (C1): $N_{G^{(1)}}(u), N_{G^{(1)}}(v) \subseteq B_{H^{(1)}}, N_{H^{(1)}}(u), N_{H^{(1)}}(v) \subseteq A_{G^{(1)}}$.
- (C2): $N_{G^{(2)}}(u), N_{G^{(2)}}(v) \subseteq B_{H^{(2)}}, N_{H^{(2)}}(u) \subseteq A_{G^{(2)}}, N_{H^{(2)}}(v) \subseteq B_{G^{(2)}}$.
- (C3): $N_{G^{(3)}}(u) \subseteq A_{H^{(3)}}, N_{G^{(3)}}(v) \subseteq B_{H^{(3)}}, N_{H^{(3)}}(u) \subseteq A_{G^{(3)}}, N_{H^{(3)}}(v) \subseteq B_{G^{(3)}}$.
- (C4): $N_{G^{(4)}}(u) \subseteq A_{H^{(4)}}, N_{G^{(4)}}(v) \subseteq B_{H^{(4)}}, N_{H^{(4)}}(u) \subseteq B_{G^{(4)}}, N_{H^{(4)}}(v) \subseteq A_{G^{(4)}}$.

Before concluding this subsection, let us present the plan for this section. In Subsection 3.2, we provide all details for constructing $G^{(\rho)}, H^{(\rho)}$. Then, in Subsection 3.3, we show that u, v are indeed $(G^{(\rho)}, H^{(\rho)})$ -exchangeable from the identity permutation. Further, Subsection 3.4 uses the lemma of Alon, Defant, and Kravitz that, for an appropriate choice of q_1, \dots, q_m , the graphs $(G^{(\rho)}, H^{(\rho)})$ are, with high probability, (q_1, \dots, q_m) -bipartite-embeddable in X, Y . Then Subsection 3.5 shows how to use (q_1, \dots, q_m) -bipartite-embeddability in order to construct embeddings ψ_G, ψ_H satisfying the conditions of Proposition 3.2. Finally, in Subsection 3.6 we complete the proof by putting together all ingredients.

3.2. Construction of $G^{(\rho)}$ and $H^{(\rho)}$. In this subsection, we describe the details of the construction of graphs $G^{(\rho)}, H^{(\rho)}$. We set the number of vertices of $G^{(\rho)}, H^{(\rho)}$ to be $m^{(\rho)} + 2$, where $m^{(1)} = m^{(2)}$ is the smallest integer larger than $(\log n)^{4/5}$ divisible by 8 and $m^{(3)} = m^{(4)} = m^{(1)} + 3$.

Throughout this subsection, we will simplify the notation $G^{(\rho)}, H^{(\rho)}, m^{(\rho)}$ and only write G, H, m , where the index ρ will be understood. As the constructions in the four cases will be essentially the same, we present them together and note where the differences arise.

Let $\ell = \lfloor \frac{1}{4}m^{1/4} \rfloor + \epsilon$, where $\epsilon \in \{0, \dots, 7\}$ is chosen so that $8|\ell$. In what follows, we will only deal with asymptotics as $m, n \rightarrow \infty$, and hence we will omit the floor symbols and ϵ .

¹Although these four cases are reminiscent of the four cases in the proof of Proposition 4.5 of [1], there is a slight difference. Namely, [1] seems to miss the subtle difference between cases (C3) and (C4) presented here, and accounts for both of them in the case (IV). Further, [1] does not make use of the symmetry $\sigma \mapsto \sigma^{-1}$ which allows us to reduce their cases (II) and (III) to a single case (C2).

As previously noted, the vertex set of G, H will be $[m] \cup \{u, v\}$. In the cases (C1) and (C2), we denote the elements of $[m]$ by

$$x_1, \dots, x_\ell, y_1, \dots, y_\ell, z_{1,1}, \dots, z_{\ell,k}, t_1, \dots, t_s, s_1, s_2, s_3, r_1, r_2, r_3, w,$$

where $k = 2\ell$ and $s = m - 7 - \ell(k + 2)$. In the cases (C3) and (C4), we denote the elements of $[m]$ by

$$x_1, \dots, x_\ell, y_1, \dots, y_\ell, z_{1,1}, \dots, z_{\ell,k}, \tilde{z}_{1,1}, \dots, \tilde{z}_{3,k}, t_1, \dots, t_s, s_1, s_2, s_3, r_1, r_2, r_3, w,$$

where $s = m - 7 - \ell(k + 2) - 3k$.

Introducing such notation for elements of $[m]$ primarily serves to distinguish the neighbors of u and v from other vertices of the graph. For example, we will have $N_G(u) = \mathcal{S} = \{s_1, s_2, s_3\}$, $N_G(v) = \mathcal{R} = \{r_1, r_2, r_3\}$, $N_H(u) = \mathcal{X} = \{x_1, \dots, x_\ell\}$, $N_H(v) = \mathcal{Y} = \{y_1, \dots, y_\ell\}$. For the sake of completeness, we also introduce notation $\mathcal{Z} = \{z_{1,1}, \dots, z_{\ell,k}\}$ in the cases (C1) and (C2), and $\mathcal{Z} = \{z_{1,1}, \dots, z_{\ell,k}, \tilde{z}_{1,1}, \dots, \tilde{z}_{3,k}\}$ in the cases (C3) and (C4). We will also assume that G, H have bipartitions $V(G) = A_G \sqcup B_G$ and $V(H) = A_H \sqcup B_H$, with $u \in A_G, A_H$ and $v \in B_G, B_H$.

We begin by describing the construction of H , since is much simpler than that of G . In the cases (C1) and (C2), we let H have a bipartiton A_H, B_H , with

$$A_H = \{u, y_1, \dots, y_\ell, w\}, \quad B_H = \{v, z_{1,1}, \dots, z_{\ell,k}, x_1, \dots, x_\ell, t_1, \dots, t_s, s_1, s_2, s_3, r_1, r_2, r_3\}$$

and the edges $\{ux_i : i \in [\ell]\}, \{vy_i : i \in [\ell]\}, \{y_i, z_{i,j} : i \in [\ell], j \in [k]\}, \{wa : a \in B_H\}$. In the cases (C3) and (C4), we let H have a bipartiton A_H, B_H , with

$$A_H = \{u, y_1, \dots, y_\ell, w, s_1, s_2, s_3\}, \quad B_H = \{v, z_{1,1}, \dots, \tilde{z}_{3,k}, x_1, \dots, x_\ell, t_1, \dots, t_s, r_1, r_2, r_3\}$$

and the edges $\{ux_i : i \in [\ell]\}, \{vy_i : i \in [\ell]\}, \{y_i, z_{i,j} : i \in [\ell], j \in [k]\}, \{wa : a \in B_H\}, \{s_i, \tilde{z}_{i,j} : i \in [3], j \in [k]\}$. Figure 4 illustrates this.

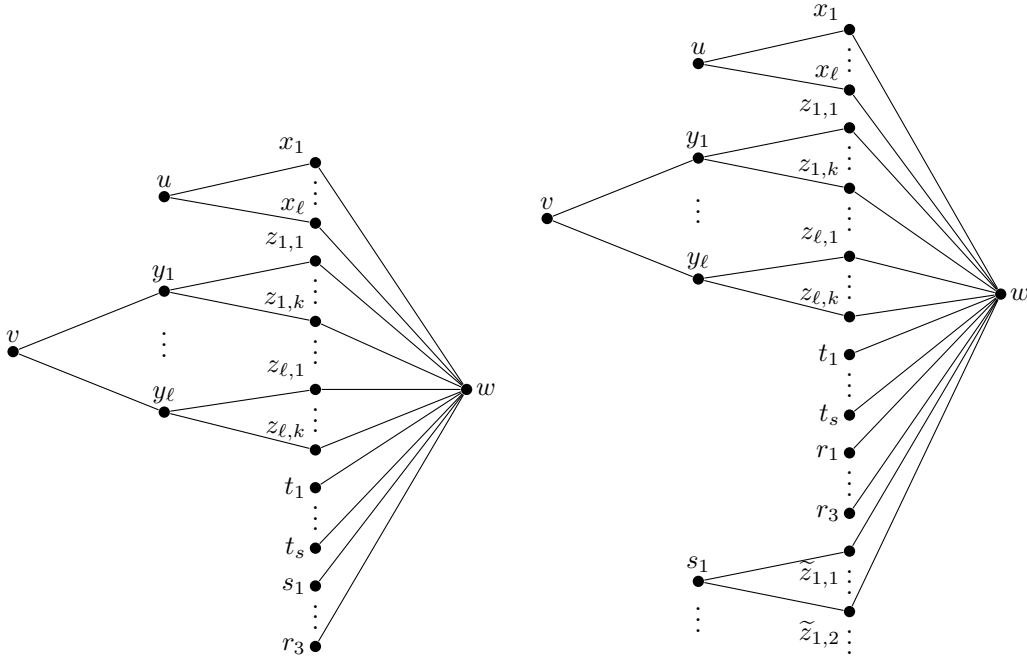


FIGURE 4. Diagrams of graphs H , for $\rho = 1, 2$ on the left and for $\rho = 3, 4$ on the right.

Now, we describe the construction of G . Set g to be the smallest integer greater than $m^{3/4}/4$ that is divisible by 4. In the cases (C1) and (C2), let G consist of a big cycle containing all vertices but u, v, y_1, \dots, y_ℓ . In the cases (C3) and (C4), on top of these, the big cycle does not contain $\tilde{z}_{1,1}, \tilde{z}_{2,1}, \tilde{z}_{3,1}$. As the big cycle is of even length, it is bipartite. We also introduce a notational convention that, if a, b are vertices of the big cycle of G , $[a, b]$ denotes clockwise interval from a to b , including all vertices between a and b as well as the endpoints. Further, we denote the clockwise distance between a and b by $d(a, b)$.

As previously noted, the vertices u, v are adjacent to s_1, s_2, s_3 and r_1, r_2, r_3 , respectively. We arrange s_1, s_2, s_3 and r_1, r_2, r_3 so that $d(s_i, r_i) = \ell' = \ell - 1$, $d(s_2, s_1) = d(s_3, s_2) = m/4$. Since $2 \nmid d(s_i, s_j)$ and $2 \nmid d(s_i, r_j)$, we conclude that s_1, s_2, s_3 lie in the same bipartite set, and that r_1, r_2, r_3 lie in the other partite set of G . Let the

partite set of G containing s_1, s_2, s_3 be B_G , and let the other partite set be A_G . Since $u \in A_G, v \in B_G$, even after adding edges us_i and vr_j , the graph G is still bipartite.

Let S_0 be the set of vertices of the big cycle that are in one of the intervals $[s_i, r_i]$, or whose neighbors on the big cycle are either s_i or r_i . For each vertex of $a \in S_0$, we will add two edges to vertices of the interval $[r_3, s_2]$. We add these edges so that the graph remains bipartite and so that all of their endpoints and s_2, r_3 have pairwise distance at least $m/(10\ell) \geq g$. The reason for adding these edges is to ensure that G remain Wilsonian even upon removing a set of vertices from the intervals $[s_i, r_i]$.

We also place the vertices x_1, \dots, x_ℓ to the interval $[r_3, s_2]$, making sure that their distance from each of the edge endpoints discussed in the previous paragraph is at least $m/(20\ell)$. Additionally, in the cases (C1), (C2) and (C3) we ensure $x_i \in A_G$ for $i \in [\ell]$ and in the case (C4) we ensure $x_i \in B_G$. Finally, we ensure that the distance between all x_i and s_2 or r_3 is at least g . Since $d(r_3, s_2) = m/4 - \ell'$, the interval $[r_3, s_2]$ is long enough for this to be done.

Similarly, we place $z_{1,1}, \dots, z_{\ell,k}, w$ (and also $\tilde{z}_{i,j}, i \in [3], j \in [2, k]$ in the cases 3 and 4) onto $[r_2, s_1]$, ensuring that the distance between any two of them is at least $m/(10\ell k)$ and distance between any of them to s_1 is at least g .

The vertices y_i will be adjacent to exactly two vertices on the big cycle — we denote them t_{2i-1} and t_{2i} . Let us set $d(t_{2i-1}, t_{2i}) = g$ and $d(t_{2i}, t_{2i+1}) = k' = k - 2$. Since both g and k are even, this means all t_i for $i \in \{1, \dots, 2\ell\}$ are in the same partite set of G , meaning that all vertices y_i are in the other partite set of G . Further, we set $d(s_1, t_1) = k - 1$ in the cases (C2), (C3), while $d(s_1, t_1) = k - 2$ in the cases (C1), (C4). These choices ensure that $\mathcal{Y} \subseteq A_G$ in the cases (C1), (C4) and that $\mathcal{Y} \subseteq B_G$ in the cases (C2), (C3) for $i \in [\ell]$.

Finally, we ensure that s_1 is adjacent to a vertex q whose distance from $t_{2\ell}$ is $k - 2$ in the cases (C1), (C4) and $k - 1$ in the cases (C2), (C3). In either case, we have $q \in A_G$, and hence G remains bipartite after adding the edge s_1q .

In the cases (C3) and (C4), the graph G has additional vertices $\tilde{z}_{1,1}, \tilde{z}_{2,1}, \tilde{z}_{3,1}, \tilde{z}_{1,2}$ outside the big cycle. Let $\tilde{t}_1, \dots, \tilde{t}_6$ be vertices of the big cycle satisfying $d(\tilde{t}_{2i-1}, s_i) = k' = k - 2$, $d(\tilde{t}_{2i}, t_{2i-1}) = g$, for $i = 1, 2, 3$. Then, we add edges $\tilde{t}_1\tilde{z}_{1,1}, \tilde{z}_{1,1}\tilde{t}_2, \tilde{t}_3\tilde{z}_{2,1}, \tilde{z}_{2,1}\tilde{t}_4, \tilde{t}_5\tilde{z}_{3,1}, \tilde{z}_{3,1}\tilde{t}_6$ to the graph G . It is simple to check that all of these edges run between different partite sets of G , which completes the description of the graphs G in all the cases. The constructions are illustrated in Figure 5.




FIGURE 5. Construction of graphs $G^{(1)}$ and $G^{(2)}$ on the left and $G^{(3)}$ and $G^{(4)}$ on the right

Since the above construction describes the graphs G fully, we will now enumerate several important properties of these graphs which will be used in the subsequent proofs. These properties should also help in understanding why the graphs G, H have been constructed in this way.

- (P1): The graph $G|_{[m]}$ has only one cycle shorter than g — this cycle contains the vertices $[s_1, t_1], y_1, [t_2, t_3], y_2, \dots, y_\ell, [t_{2\ell}, q], s_1$ and has length $k \cdot (l + 1) - 2 \geq lk$.

- (P2): The smallest distance, in the graph G between any two elements of the set $\mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z} \cup \mathcal{S} \cup \mathcal{R}$ is at least $\ell' = \ell - 1$.
- (P3): Let $P = \{y_1, \dots, y_\ell, u, v\} \cup \bigcup_{i=1}^3 [s_i, r_i]$ in the cases (C1) and (C2) and $P = \{y_1, \dots, y_\ell, u, v\} \cup \bigcup_{i=1}^3 [s_i, r_i] \cup \{\tilde{z}_{1,1}, \tilde{z}_{2,1}, \tilde{z}_{3,1}\}$ in the cases (C3) and (C4). If S_1 is a subset of P , we note that $G|_{V(G) \setminus S_1}$ is Wilsonian. In other words, removing a vertices of P from G does not change the fact G is Wilsonian.
- (P4): The number of edges of G is $|E(G)| \leq |V(G)| + 4\ell + 15 \leq |V(G)| + 5\ell$. Moreover, for any $J \subseteq V(G)$, the subgraph $G|_J$ has at most $|J| + 5\ell$ edges.

3.3. Showing that u and v are $(G^{(\rho)}, H^{(\rho)})$ -exchangeable from the identity permutation. Now, we will prove the most important property of graphs $(G^{(\rho)}, H^{(\rho)})$ — that the vertices u and v are $(G^{(\rho)}, H^{(\rho)})$ -exchangeable starting from the identity permutation.

Lemma 3.3. *For $\rho \in \{1, 2\}$, the vertices u and v are $(G^{(\rho)}, H^{(\rho)})$ -exchangeable from the identity permutation.*

Proof. Again, for simplicity of notation, we omit the index (ρ) in this proof. The general strategy of the proof of this lemma hinges on repeated applications of Wilson's theorem, exploiting the fact that $H|_{\{w\} \cup B_H - \{v\}}$ is a star centered at w .

Let us fix two vertices a_1 and a_2 in the same partite set of G , and let $\tau : V(G) \rightarrow V(H)$ be an arbitrary arrangement. If $\tau' : V(G) \rightarrow V(H)$ is any arrangement differing from τ only on elements of $\tau^{-1}(\{w\} \cup B_H - \{v\})$, and if the graph $G|_{\tau^{-1}(\{w\} \cup B_H - \{v\})}$ is Wilsonian and contains a_1, a_2 , then τ is in the same connected component as either τ' or $\tau' \circ (a_1 a_2)$. To see why this is true, note that the arrangements τ' and $\tau' \circ (a_1 a_2)$ are not in the same connected component, due to Proposition 2.4. Since Wilson's theorem implies the graph $\text{FS}(G|_{\tau^{-1}(\{w\} \cup B_H - \{v\})}, H|_{\{w\} \cup B_H - \{v\}})$ has exactly two connected components, one of τ' and $\tau' \circ (a_1 a_2)$ must be in the same connected component as τ .

We will the above apply simple fact with a_1, a_2 being vertices satisfying $d(a_1, q) = 4$ and $d(a_2, q) = 2$. The fact that $G|_{\tau^{-1}(\{w\} \cup B_H - \{v\})}$ is Wilsonian will follow from $\tau^{-1}(A_H - \{w\} \cup \{v\}) \subseteq P$ and property (P3) of the graph G . Hence, whenever we can ensure that the vertices of $\tau^{-1}(A_H - \{w\} \cup \{v\})$ all lie in P , we can obtain any arrangement of the remaining vertices, up to possibly transposing a_1 and a_2 (we will omit this clause in the future, but it will be understood whenever we apply Wilson's theorem).

We will think of arrangements $\tau : V(G) \rightarrow V(H)$ as placing the labels corresponding to vertices of H onto the vertices of G so that the label $p \in V(H)$ is placed onto the vertex $\tau^{-1}(p) \in V(G)$. Moreover, when we refer to the (G, H) -friendly swap ab , it will be understood the swap involves labels $a, b \in V(H)$. Finally, the reader is encouraged to draw out the following steps which transform Id into $(u v)$.

Let τ_1 be an arrangement in which $\tau_1^{-1}(z_{1,1}), \dots, \tau_1^{-1}(z_{1,k-1})$ appear consecutively in clockwise order on $[s_1, t_1]$. By Wilson's theorem, there exists a sequence of (G, H) -friendly swaps transforming Id to τ_1 .

Let τ_2 be an arrangement obtained from τ_1 by performing swaps $y_1 z_{1,k-1}, y_1 z_{1,k-2}, \dots, y_1 z_{1,1}$. Note that τ_2 satisfies $\tau_2^{-1}(y_1) = s_1$, and that it can be obtained from Id by a sequence of (G, H) -friendly swaps. In essence, we can think of this sequence as moving the label $y_1 \in V(H)$ from the vertex $y_1 \in V(G)$ to the vertex $s_1 \in V(G)$.

Let τ_3 be the arrangement which differs from τ_2 only on values $\tau_2^{-1}(\{z_{2,1}, \dots, z_{2,k-1}\}) \cup \{y_1\} \cup [t_2, t_3]$, and in which the vertices $\tau_3^{-1}(z_{2,2}), \dots, \tau_3^{-1}(z_{2,k-1})$ appear consecutively in clockwise order on $[t_2, t_3]$ and which satisfies $\tau_3^{-1}(z_{2,1}) = y_1$. By Wilson's theorem, there exists a sequence of (G, H) -friendly swaps transforming τ_2 to τ_3 .

Let τ_4 be an arrangement obtained from τ_3 by performing swaps $y_2 z_{2,k-1}, y_2 z_{2,k-2}, \dots, y_2 z_{2,1}$. Note that τ_4 satisfies $\tau_4^{-1}(y_2) = y_1$. Intuitively, τ_4 is obtained from τ_2 by performing a sequence of swaps which moves the label $y_2 \in V(H)$ from the vertex $y_2 \in V(G)$ to the vertex $y_1 \in V(G)$.

Repeatedly applying sequences of swaps analogous to those taking y_1 to s_1 or y_2 to y_1 , one can arrive at an arrangement τ_5 in which vertices $\tau_5^{-1}(y_1), \dots, \tau_5^{-1}(y_\ell)$ appear consecutively in clockwise order on $[s_1, r_1]$.

Let τ_6 be the arrangement obtained from τ_5 by performing the swaps $v y_\ell, \dots, v y_1$ which satisfies $\tau_6^{-1}(v) = s_1$ and $\tau_6^{-1}(y_\ell) = v$.

Let τ_7 be the arrangement which differs from τ_6 only on values $\tau_6^{-1}(\{x_1, \dots, x_\ell\}) \cup [s_3, r_3]$, and in which the vertices $\tau_6^{-1}(x_1), \dots, \tau_6^{-1}(x_\ell)$ appear consecutively in clockwise order on $[s_3, r_3]$. By Wilson's theorem, there exists a sequence of (G, H) -friendly swaps transforming τ_6 to τ_7 .

Let τ_8 be the arrangement obtained from τ_7 by performing the swaps $u x_1, \dots, u x_{\ell-1}$ which satisfies $\tau_8^{-1}(x_1) = u$ and $\tau_8^{-1}(u)$ is adjacent to the vertex $r_3 \in V(G)$.

Let $\tau_9 = (z_{\ell,1} x_1) \circ \tau_8$. By Wilson's theorem, there exists a sequence of (G, H) -friendly swaps transforming τ_8 to τ_9 . The arrangement τ_9 satisfies $\tau_9^{-1}(z_{\ell,1}) = u$.

Let τ_{10} be the arrangement arising from τ_9 after performing swaps $y_\ell z_{\ell,1}, y_\ell v$. In τ_{10} , the vertices $\tau_{10}^{-1}(y_\ell), \tau_{10}^{-1}(y_1), \tau_{10}^{-1}(y_2), \dots, \tau_{10}^{-1}(y_{\ell-1})$ appear consecutively in clockwise order on $[s_1, r_1]$. Also, $\tau_{10}^{-1}(v) = u$.

Let τ_{11} be the arrangement which differs from τ_{10} only on values $\tau_{10}^{-1}(\{x_1, x_2\}) \cup \{r_3, v\}$, and which satisfies $\tau_{11}^{-1}(x_1) = r_3, \tau_{11}^{-1}(x_2) = v$. By Wilson's theorem, there exists a sequence of (G, H) -friendly swaps transforming τ_{10} to τ_{11} .

Let τ_{12} be the arrangement obtained from τ_7 by performing the swap ux_1 , satisfying $\tau_{12}(u) = v$ and $\tau_{12}(v) = u$. It remains to show that all other labels can be returned to their original positions. By applying the reverse of the process which brought the labels $y_1, \dots, y_{\ell-1}$ onto the vertices $[s_1, r_1]$, it is simple to see that there exists an arrangement τ_{13} which satisfies $\tau_{13}|_{\{u,v\}} = \tau_{12}|_{\{u,v\}}$ and $\tau_{13}(y_i) = y_i$ for $i \in [l-1]$.

Let τ_{14} be the arrangement which differs from τ_{13} only on values $\tau_{13}^{-1}(\{z_{\ell,1}, \dots, z_{\ell,k-1}\}) \cup [t_{2\ell}, q] \cup \{y_\ell\}$, and in which vertices $\tau_{14}^{-1}(z_{\ell,k-1}), \dots, \tau_{14}^{-1}(z_{\ell,2})$ appear consecutively in clockwise order on $[t_{2\ell}, q]$ and $\tau_{14}^{-1}(z_{\ell,1}) = y_\ell$. By Wilson's theorem, there exists a sequence of (G, H) -friendly swaps transforming τ_{13} to τ_{14} .

Let τ_{15} be the arrangement obtained from τ_{14} by performing the swaps $y_\ell z_{\ell,k-1}, \dots, y_\ell z_{\ell,1}$ which satisfies $\tau_{15}^{-1}(y_\ell) = y_\ell$.

Finally, note that τ_{15}^{-1} fixes all elements of $A_H - \{w\} \cup \{v\}$, except u and v which are swapped. Hence, using Wilson's theorem, we conclude that there exists a sequence of (G, H) -friendly swaps transforming τ_{14} to $(u v)$. However, note that we actually obtain $\tau_{15} = (u v)$ only up to possibly transposing a_1 and a_2 , even though we did not explicitly state this after every application of Wilson's theorem in the proof.

Hence, we either have $\tau_{15} = (u v)$ or $\tau_{15} = (u v) \circ (a_1 a_2)$. However, note that it is not possible that we obtained the latter arrangement, starting from the identity permutation, due to parity constraints. Therefore, we conclude $\tau_{15} = (u v)$ and the proof is complete. \square

Lemma 3.4. *For $\rho \in \{3, 4\}$, the vertices u and v are $(G^{(\rho)}, H^{(\rho)})$ -exchangeable from the identity permutation.*

Proof. In the cases (C3) and (C4), the identity permutation satisfies $A_H - \{w\} \cup \{v\} \subseteq P$, and hence there exists a sequence of $(G^{(\rho)}, H^{(\rho)})$ -friendly swaps taking Id to τ_0 where τ_0 has the property that $\tau_0^{-1}(\tilde{z}_{1,1}), \dots, \tau_0^{-1}(\tilde{z}_{1,k-2})$ are consecutively arranged in counterclockwise order on the big cycle, $\tau_0^{-1}(\tilde{z}_{1,1})$ being the neighbor of s_1 and $\tau_0^{-1}(\tilde{z}_{1,k-2}) = \tilde{t}_1$. Moreover, assume $\tau_0^{-1}(\tilde{z}_{1,k-1}) = \tilde{z}_{1,1}$. Using Wilson's theorem as described above, we conclude that τ_0 can be obtained from Id by a sequence of $(G^{(\rho)}, H^{(\rho)})$ -friendly swaps. If we consider the sequence of swaps $s_1 \tilde{z}_{1,1}, s_1 \tilde{z}_{1,2}, \dots, s_1 \tilde{z}_{1,k-1}$ applied on τ_0 , we obtain an arrangement τ_1 having $\tau_1(s_1) = \tilde{z}_{1,1}$ and $\tau_1(\tilde{z}_{1,1}) = s_1$. Finally, using one more application of Wilson's theorem, there exists a sequence of $(G^{(\rho)}, H^{(\rho)})$ -friendly swaps which transforms τ_1 into an arrangement $\tau_2 = (s_1 \tilde{z}_{1,1})$. Applying the same argument to s_2, s_3 , we obtain an arrangement τ_3 which transposes elements s_i and $\tilde{z}_{i,1}$, and leaves all other elements fixed. Since the elements u and v were not involved in any of the swaps transforming Id to τ_3 , it is not hard to see that the same sequence of swaps transforms $(u v)$ to $\tau_3 \circ (u v)$.

The important thing to note now is that eliminating s_1, s_2, s_3 from $H^{(\rho)}$ and $\tilde{z}_{1,1}, \tilde{z}_{2,1}, \tilde{z}_{3,1}$ from $G^{(\rho)}$ essentially reduces to the case $\rho \in \{1, 2\}$. Let us now explain why more formally.

Note that $G^{(\rho)} - \{\tilde{z}_{1,1}, \tilde{z}_{2,1}, \tilde{z}_{3,1}\}$ for $\rho \in \{3, 4\}$ is isomorphic, as an unlabeled graph, to $G^{(5-\rho)}$, and that $H^{(\rho)} - \{s_1, s_2, s_3\}$ for $\rho \in \{3, 4\}$ is isomorphic, again in the unlabeled sense, to $H^{(5-\rho)}$. In other words, there exist embeddings $\varphi_G : G^{(5-\rho)} \rightarrow G^{(\rho)}$ and $\varphi_H : H^{(5-\rho)} \rightarrow H^{(\rho)}$ which map between the arrangement Id on $(G^{(5-\rho)}, H^{(5-\rho)})$ and arrangement τ_3 on $(G^{(\rho)}, H^{(\rho)})$. Alternatively, we may say that $\tau_3 \circ \varphi_G = \varphi_H \circ \text{Id}$ and $\varphi_H(u) = u, \varphi_H(v) = v$. Since Lemma 3.3 shows that u and v are $(G^{(5-\rho)}, H^{(5-\rho)})$ -exchangeable from Id , using Proposition 3.2 translates to u and v being $(G^{(\rho)}, H^{(\rho)})$ -exchangeable from τ_3 . Combining the sequences obtaining transformations $\text{Id} \rightarrow \tau_3, \tau_3 \rightarrow \tau_3 \circ (u v)$ and $\tau_3 \circ (u v) \rightarrow (u v)$, we conclude that Id and $(u v)$ are connected in the friends-and-strangers graph $\text{FS}(G^{(\rho)}, H^{(\rho)})$, which implies u, v are $(G^{(\rho)}, H^{(\rho)})$ -exchangeable from Id . \square

3.4. Showing that $G^{(\rho)}, H^{(\rho)}$ are (q_1, \dots, q_m) -bipartite-embeddable in X, Y . In this section, we will prove that, with high probability, the graphs $G^{(\rho)}|_{[m]}, H^{(\rho)}|_{[m]}$ are (q_1, \dots, q_m) -bipartite-embeddable in (X, Y) with high probability. The main technical tool which allows us to show this is Lemma 4.2 from [1], which states the following.

Lemma 3.5 ([1]). *Let G' and H' be bipartite graphs on the vertex set $[m]$ with bipartitions $V(G') = A_{G'} \sqcup B_{G'}$, $V(H') = A_{H'} \sqcup B_{H'}$. Let n, q_1, \dots, q_m be positive integers such that $Q := q_1 + \dots + q_m \leq 2n$. For every set $J \subseteq [m]$, let $\beta(J) = |E(G'|_J)| + |E(H'|_J)|$. Choose $p \in [0, 1]$, and let X, Y be independently chosen random bipartite graphs in $\mathcal{G}(K_{n,n}, p)$. Suppose that for every $J \subseteq [m]$ with $\beta(J) \geq 1$ we have*

$$(2) \quad p^{\beta(J)} \prod_{j \in J} q_j \geq 3 \cdot 2^{m+1} Q \log(2n).$$

Then, the probability that the pair (G', H') is (q_1, \dots, q_m) -bipartite-embeddable in (X, Y) is at least $1 - (2n)^{-Q}$.

In Lemma 3.6, we check that condition (2) indeed holds for $G^{(\rho)}|_{[m]}, H^{(\rho)}|_{[m]}$ under the appropriate choice of values q_1, \dots, q_m . Let $\Gamma = N_H(u) \cup N_H(v) \cup N_G(u) \cup N_G(v) = \mathcal{X} \cup \mathcal{Y} \cup \mathcal{S} \cup \mathcal{R}$.

Lemma 3.6. *Let X, Y be random bipartite graphs, independently chosen from $\mathcal{G}(K_{n,n}, p)$, where p satisfies (1). Moreover, suppose that $q_i = \frac{n}{2m}$ for $i \in [m] - \Gamma$ and $q_i = \frac{pn}{3\ell}$ for $i \in \Gamma$. Then, for any $\rho \in \{1, 2, 3, 4\}$, the pair of graphs $(G^{(\rho)}|_{[m]}, H^{(\rho)}|_{[m]})$ is (q_1, \dots, q_m) -bipartite-embeddable in (X, Y) with high probability.*

Proof. This proof is inspired by the proof of Lemma 3.4 in [1], with the main difference coming only from the complexity of our construction. To simplify notation, let $G' = G^{(\rho)}|_{[m]}, H' = H^{(\rho)}|_{[m]}$, and $Q = \sum_{i=1}^m q_i$. Using Lemma 3.5, the statement about embeddability of (G', H') reduces to checking that

$$p^{\beta(J)} \prod_{j \in J} q_j \geq 3 \cdot 2^{m+1} Q \log(2n),$$

for all J with $\beta(J) \geq 1$. Letting $\gamma(J) = |\Gamma \cap J|$ and using $Q \leq n$, the above inequality reduces to

$$(3) \quad p^{\beta(J)} \left(\frac{pn}{3\ell}\right)^{\gamma(J)} \left(\frac{n}{2m}\right)^{|J|-\gamma(J)} \geq 3 \cdot 2^{m+1} n \log 2n.$$

Since $2m > 3\ell$, one can bound the left hand side of equation (3) in the following way:

$$p^{\beta(J)} \left(\frac{pn}{3\ell}\right)^{\gamma(J)} \left(\frac{n}{2m}\right)^{|J|-\gamma(J)} \geq p^{\beta(J)+\gamma(J)-2|J|} \left(\frac{p^2n}{2m}\right)^{|J|} \left(\frac{2m}{3\ell}\right)^{\gamma(J)} \geq p^{\beta(J)+\gamma(J)-2|J|} \left(\frac{p^2n}{2m}\right)^{|J|}$$

We will now bound $\beta(J)$ and $\gamma(J)$ by analyzing the structure of the graphs G' and H' . Recall that H' has a bipartition $V(H') = A_{H'} \sqcup B_{H'}$, where $A_{H'} = A_H - \{u\}$ and $B_{H'} = B_H - \{v\}$. With this in mind, we can express the number of edges of $H'|_J$ as the sum of degrees of vertices in one of the partite classes, $|E(H'|_J)| = \sum_{a \in J \cap A_{H'}} |N_{H'}(a) \cap J|$. Expanding out this expression gives:

$$(4) \quad |E(H'|_J)| = \mathbf{1}_{w \in J} |N_{H'}(w) \cap J| + \sum_{a \in A_{H'} - \{w\}} |N_{H'}(a) \cap J| \leq \mathbf{1}_{w \in J} |J \cap B_{H'}| + |J \cap \mathcal{Z}|,$$

where $\mathbf{1}_{w \in J}$ denotes the indicator variable which equals 1 when $w \in J$ and 0 otherwise.

Let $\Gamma' = (\mathcal{X} \cup \mathcal{Z} \cup \mathcal{S} \cup \mathcal{R}) \cap B_{H'}$, and define $\gamma'(J) = |J \cap \Gamma'|$ and $\alpha(J) = |E(G'|_J)|$. Then, the bound (4) applied to give:

$$(5) \quad \begin{aligned} \beta(J) + \gamma(J) - 2|J| &= (|E(G'|_J)| + |E(H'|_J)|) + (\gamma'(J) + |J \cap (A_{H'} - \{w\})| - |J \cap \mathcal{Z}|) - 2|J| \\ &\leq \alpha(J) + \mathbf{1}_{w \in J} |J \cap B_{H'}| + \gamma'(J) + |J \cap (A_{H'} - \{w\})| - 2|J| \\ &\leq \alpha(J) + \gamma'(J) - |J| - \mathbf{1}_{w \in J}. \end{aligned}$$

We now consider five cases and apply the above bounds in resolving them.

Case 1: Suppose that $|J| \geq g$. Since $|E(G'|_J)| \leq |V(G')| + 5\ell$, by the property (P4) of the graphs G' and since G' is a connected graph, we also have $\alpha(J) = |E(G'|_J)| \leq |J| + 5\ell$. Therefore, one can apply the bound (5) and show (3) in the following way:

$$\begin{aligned} p^{\beta(J)+\gamma(J)-2|J|} \left(\frac{p^2n}{2m}\right)^{|J|} &\geq p^{\alpha(J)+\gamma'(J)-|J|} \left(\frac{p^2n}{2m}\right)^{|J|} \geq p^{5\ell+|\Gamma'|} \left(\frac{\exp(20(\log n)^{4/5})}{2m}\right)^g \\ &\geq n^{-(6\ell+k(\ell+3)+6)/2} \exp\left(5(\log n)^{4/5} m^{3/4} - g \log(2m)\right) \\ &\geq \exp\left(5(\log n)^{7/5} - \frac{3}{2}\ell^2 \log n\right) \geq \exp\left(3(\log n)^{7/5}\right) \gg 3 \cdot 2^{m+1} n \log(2n). \end{aligned}$$

Case 2: Suppose that $|J| < g$ and $\gamma'(J) > c(J)$, where $c(J)$ is the number of connected components of J . Since there exists exactly one cycle of length less than g in G , we must have $\alpha(J) \leq |J| - c(J) + 1$. On the other hand, since the distance between any two elements of Γ' is at least $\ell' = \ell - 1$, we have $|J| \geq (\gamma'(J) - c(J))\ell'$. Therefore

$$\begin{aligned} p^{\beta(J)+\gamma(J)-2|J|} \left(\frac{p^2n}{2m}\right)^{|J|} &\geq p^{\alpha(J)+\gamma'(J)-|J|} \left(\frac{p^2n}{2m}\right)^{|J|} \geq p^{\gamma'(J)-c(J)+1} \left(\frac{\exp(20 \log n^{4/5})}{2m}\right)^{(\gamma'(J)-c(J))\ell'} \\ &\geq n^{-1/2} \left(n^{-1/2} \exp(4 \log n)\right)^{\gamma'(J)-c(J)} \geq n^3 \gg 3 \cdot 2^{m+1} n \log(2n). \end{aligned}$$

Case 3: Suppose that $\gamma'(J) \leq c(J)$ and $\ell \leq |J| < g$. As discussed in the previous case, we have $\alpha(J) \leq |J| - c(J) + 1$, which implies

$$\begin{aligned} p^{\beta(J) + \gamma(J) - 2|J|} \left(\frac{p^2 n}{2m} \right)^{|J|} &\geq p^{\alpha(J) + \gamma'(J) - |J|} \left(\frac{p^2 n}{2m} \right)^{|J|} \geq p^{\gamma'(J) - c(J) + 1} \left(\frac{\exp(20(\log n)^{4/5})}{2m} \right)^\ell \\ &\geq n^{-1/2} \exp\left(5 \log n - (\log n)^{1/5} \log(2m)\right) \geq n^4 \gg 3 \cdot 2^{m+1} n \log(2n). \end{aligned}$$

Case 4: Suppose that $\gamma'(J) \leq c(J)$, $|J| < \ell$ and $w \in J$. Since $|J|$ is smaller than the girth of G' , $G'|_J$ must be a forest and therefore $\alpha(J) = |J| - c(J)$. Also, as $\beta(J) \geq 1$, J must induce at least one edge in G' or H' , and hence $|J| \geq 2$. Finally, since elements of $(\Gamma' \cup \{w\})$ are at mutual distance at least ℓ' , we conclude that every connected component of $G'|_J$ contains at most one element of $\Gamma' \cup \{w\}$, implying $\gamma'(J) + 1 = |J \cap (\Gamma' \cup \{w\})| \leq c(J)$. Combining these facts, one has:

$$\begin{aligned} p^{\beta(J) + \gamma(J) - 2|J|} \left(\frac{p^2 n}{2m} \right)^{|J|} &\geq p^{\alpha(J) + \gamma'(J) - |J| - 1} \left(\frac{p^2 n}{2m} \right)^{|J|} \geq p^{\gamma'(J) - c(J) - 1} \left(\frac{p^2 n}{2m} \right)^{|J|} \\ &\geq p^{-2} \left(\frac{p^2 n}{2m} \right)^2 = \frac{p^2 n^2}{(2m)^2} \geq \frac{\exp(20(\log n)^{4/5}) n}{4m^2} \gg 3 \cdot 2^{m+1} n \log 2n. \end{aligned}$$

Case 5: Suppose that $\gamma'(J) \leq c(J)$, $|J| < \ell$, and $w \notin J$. As in the previous case, we infer that J is a forest, and that $\alpha(J) = |J| - c(J)$. If J induces no edges on H' we have

$$\beta(J) + \gamma(J) - 2|J| = \alpha(J) + \gamma(J) - 2|J| = \gamma(J) - c(J) - |J|.$$

Note that every two elements of Γ are at distance at least $\ell' = \ell - 1$, which gives $\gamma(J) \leq c(J)$, and hence

$$\begin{aligned} p^{\beta(J) + \gamma(J) - 2|J|} \left(\frac{p^2 n}{2m} \right)^{|J|} &\geq p^{-2} \left(\frac{p^2 n}{2m} \right)^2 = \frac{p^2 n^2}{(2m)^2} \\ &\geq \frac{\exp(20(\log n)^{4/5}) n}{4m^2} \gg 3 \cdot 2^{m+1} n \log 2n. \end{aligned}$$

In the case that J induces an edge in H' one can use the fact that $w \notin J$ to improve the bound (4)

$$\begin{aligned} \beta(J) + \gamma(J) - 2|J| &= \alpha(J) + \gamma(J) + |J \cap \mathcal{Z}| - 2|J| \\ &= \gamma'(J) - c(J) + |J \cap A_{H'}| - |J| = \gamma'(J) - c(J) - |J \cap B_{H'}|. \end{aligned}$$

Since $|J| < \ell$ and the distance between any two vertices of $(A_{H'} - \{w\}) \cup \mathcal{Z}$ is at least $\ell' = \ell - 1$, we conclude that $c(J) \geq |J \cap (A_{H'} \cup \mathcal{Z})|$. Noting that J intersects both $A_{H'}$ and $B_{H'}$, we see that $c(J) \geq \gamma'(J) + 1$ and $|J \cap B_{H'}| \geq 1$. Hence, $\gamma'(J) - c(J) - |J \cap B_{H'}| \leq -2$ and so

$$\begin{aligned} p^{\beta(J) + \gamma(J) - 2|J|} \left(\frac{p^2 n}{2m} \right)^{|J|} &\geq p^{-2} \left(\frac{p^2 n}{2m} \right)^2 = \frac{p^2 n^2}{(2m)^2} \\ &\geq \frac{\exp(20(\log n)^{4/5}) n}{4m^2} \gg 3 \cdot 2^{m+1} n \log 2n, \end{aligned}$$

which completes the proof of Lemma 3.6. \square

3.5. Constructing ψ_G and ψ_H . Before showing how to construct ψ_G, ψ_H and completing the proof, we have one additional step. Namely, in the construction of ψ_G, ψ_H , we will need that the arrangement σ sends approximately the same number of vertices from A_X to A_Y and B_Y . More precisely, we say that an arrangement $\sigma : V(X) \rightarrow V(Y)$ is *balanced* if $2n/3 \geq |\sigma(A_X) \cap A_Y| \geq n/3$. Otherwise, σ is *unbalanced*. Although it may seem at first that this definition is not symmetric in A_X and B_X , the symmetry comes from the fact that $|\sigma(A_X) \cap A_Y| = n - |\sigma(B_X) \cap A_Y| = |\sigma(B_X) \cap B_Y| = n - |\sigma(A_X) \cap B_Y|$. Together with Proposition 2.2, the following result essentially shows that we may, without loss of generality, assume that we are exchanging u_0 and v_0 starting from a balanced arrangement. Although Proposition 3.7 is stated as a result about exchangeability in various arrangements, the underlying argument is really about large matchings in random bipartite graphs.

Proposition 3.7. *Let X, Y be random bipartite graphs, independently chosen from $\mathcal{G}(K_{n,n}, p)$. With high probability, the following statement holds: for any unbalanced arrangement $\sigma : V(X) \rightarrow V(Y)$ and any vertices $u_0, v_0 \in V(Y)$, there exists a sequence of (X, Y) -friendly swaps, not involving the vertices u_0 and v_0 , which transforms the arrangement σ into a balanced arrangement $\sigma' : V(X) \rightarrow V(Y)$.*

Proof. The first step of the proof will be to show that for a fixed permutation $\sigma : V(X) \rightarrow V(Y)$, the required sequence exists with probability at least $1 - e^{-\Omega(p^2 n^2)}$. After showing this, a simple application of the union bound will suffice to complete the proof.

Having fixed an unbalanced arrangement $\sigma : V(X) \rightarrow V(Y)$, we will consider the case when $|\sigma(A_X) \cap A_Y| > 2n/3$. Note that the other case, when $|\sigma(A_X) \cap A_Y| < n/3$, can be reduced to the first one by exchanging labels A_Y and B_Y . The main idea of the proof will be to show that, with high probability, there exist many pairs of vertices $a \in A_Y, b \in B_Y$ with $\sigma(a) \in A_X, \sigma(b) \in B_X$ which can be swapped using a (X, Y) -friendly swap, since performing roughly $n/3$ such swaps will transform σ into a balanced arrangement σ' . We will then show that choosing many such disjoint pairs (a, b) corresponds to finding a large matching in a certain random bipartite graph, which will in turn be accomplished using König's theorem.

Before passing to the proof, let us introduce some additional notation. Define $C_Y = \sigma(A_X) \cap A_Y$ and $D_Y = \sigma(B_X) \cap B_Y$, and let C_X, D_X be the preimages of these sets under σ , $C_X = \sigma^{-1}(C_Y), D_X = \sigma^{-1}(D_Y)$. Also, let us denote the size of each of the sets C_X, D_X, C_Y, D_Y by n' , where $n' \geq 2n/3$.

For a fixed permutation $\sigma : V(X) \rightarrow V(Y)$ and random $X, Y \sim \mathcal{G}(n, p)$, we construct a bipartite graph Z with the bipartition $V(Z) = C_Y \sqcup D_Y$ such that $cd \in E(Z)$ if and only if $cd \in E(Y)$ and $\sigma^{-1}(c)\sigma^{-1}(d) \in E(X)$. Since σ is fixed, the events $cd \in E(Z)$ are mutually independent for all $c \in C_Y, d \in D_Y$, and they occur with probability p^2 . Hence, Z has a distribution of a random bipartite graph $\mathcal{G}(K_{n', n'}, p^2)$.

Suppose that Z has a matching $M = \{c_1 d_1, \dots, c_k d_k\}$ of size $k > n/2$. Assuming $u_0, v_0 \notin \{c_1, d_1, \dots, c_{n/3}, d_{n/3}\}$, we see that performing the (X, Y) -friendly swaps $c_1 d_1, \dots, c_{n/3} d_{n/3}$ gives a sequence of swaps which transforms σ into a balanced arrangement σ' . Hence, it suffices to show that the size of the maximum matching of Z , denoted by $\nu(Z)$, is at least $n/2$ with probability at least $1 - e^{-\Omega(p^2 n^2)}$.

By König's theorem, the size of the largest matching in Z is equal to the size of the smallest vertex cover of Z , which is denoted by $\tau(Z)$. Therefore, we will begin by bounding the probability that a fixed set V , of size $n/2$, is a vertex cover of Z . Note that a set V covers at most $|V|n'$ edges of the graph $K_{n', n'}$. Hence, the probability that V is a vertex cover of Z is equal to the probability that none of the remaining $n'(n' - |V|)$ edges exist in $E(Z)$. In other words, we have;

$$\mathbb{P}[V \text{ is a vertex cover of } Z] \leq (1 - p^2)^{n'(n' - |V|)} \leq e^{-\Omega(p^2 n^2)}.$$

Hence, the probability that there exists a vertex cover of Z of size $\leq n/2$ is at most

$$\mathbb{P}[\nu(Z) \leq n/2] \leq \binom{n}{n/2} e^{-\Omega(p^2 n^2)} \leq 2^n e^{-\Omega(p^2 n^2)} \leq e^{-\Omega(p^2 n^2)}.$$

The above computation shows that the probability that the required sequence exists for a fixed choice of σ is at least $1 - e^{-\Omega(p^2 n^2)}$. Hence, a union bound over all choices of σ shows that the statement of the proposition fails with probability $\leq n! e^{-\Omega(p^2 n^2)} = o(1)$, as $n \rightarrow \infty$. This completes the proof of the Proposition 3.7. \square

The last prerequisite we need for the proof is the way to construct ψ_G and ψ_H when $G^{(\rho)}|_{[m]}, H^{(\rho)}|_{[m]}$ are (q_1, \dots, q_m) -bipartite-embeddable in X, Y .

Proposition 3.8. *Let X, Y be random bipartite graphs, independently chosen from $\mathcal{G}(K_{n, n}, p)$. With high probability, the following statement holds: For any two vertices $u_0, v_0 \in V(Y)$ satisfying $u_0 v_0 \in E(Y), \sigma^{-1}(u_0) \in A_X, \sigma^{-1}(v_0) \in A_Y$ and every balanced $\sigma : V(X) \rightarrow V(Y)$ satisfying the assumptions of the case $(C\rho)$, for $\rho \in \{1, 2, 3, 4\}$, there exist embeddings $\psi_G : V(G^{(\rho)}) \rightarrow V(X)$ and $\psi_H : V(H^{(\rho)}) \rightarrow V(Y)$ with $\sigma \circ \psi_G = \psi_H \circ \text{Id}$ and $\psi_H(u) = u_0, \psi_H(v) = v_0$.*

Proof. The main idea of the proof is to carefully choose the sets V_1, \dots, V_m such that the adjacency of the corresponding vertices of X, Y with u_0, v_0, u'_0, v'_0 is guaranteed by the choice of V_1, \dots, V_m . For the sake of concreteness, we will describe in detail how these sets are chosen only in the case $\rho = 1$, as the choices are made in a completely analogous way in all other cases.

From the condition $\sigma \circ \psi_G = \psi_H \circ \text{Id}$ and $\psi_H(u) = u_0, \psi_H(v) = v_0$, we must have $\psi_G(u) = \sigma^{-1}(u_0) = u'_0$ and $\psi_G(v) = v'_0$.

Since $p \gg \frac{\log n}{n}$, we have that, with high probability, all vertices of X, Y have degrees at least $2pn/3$, which we assume to be the case. For $i \in N_{G^{(1)}}(u)$, we choose V_i of size $\frac{pn}{3\ell}$ to lie in $\sigma(N_X(u'_0)) \cap B_Y$, and for $i \in N_{G^{(1)}}(v)$, we choose V_i of size $\frac{pn}{3\ell}$ to lie in $\sigma(N_X(v'_0)) \cap B_Y$. By assumption σ majority-maps $N_X(u'_0)$ and $N_X(v'_0)$ to B_Y , so the sizes of the intersections $\sigma(N_X(u'_0)) \cap B_Y$ and $\sigma(N_X(v'_0)) \cap B_Y$ are at least $\frac{pn}{3}$. Hence, one can indeed choose l disjoint sets V_i of size $\frac{pn}{3\ell}$ satisfying the required properties.

Similarly, for $i \in N_{H^{(1)}}(u)$ we choose a set $V_i \subset N_Y(u_0) \cap \sigma(A_X)$ of size $\frac{pn}{3\ell}$ and for $i \in N_{H^{(1)}}(v)$ we choose $V_i \subset N_Y(v) \cap \sigma(A_X)$, again of size $\frac{pn}{3\ell}$. For vertices $i \in [m]$ which are not neighbors of u, v in either $G^{(1)}$ or $H^{(1)}$, say

for $i \in A_G \cap A_H - \Gamma$, we choose a set $V_i \subset \sigma(A_X) \cap A_Y$, of size $\frac{n}{4m}$. Since σ is balanced, we have $|\sigma(A_X) \cap A_Y| \geq n/3$ and hence the sets V_i for $i \in A_G \cap A_H - \Gamma$ can be chosen to be disjoint from all previously chosen sets. The sets corresponding to vertices $i \in A_G \cap B_H, B_G \cap A_H, B_G \cap B_H$ are chosen analogously.

From our choice of V_1, \dots, V_m it is clear that this list is admissible, and that $|V_i| = q_i$, where q_i is defined in the statement of Lemma 3.6. Hence, by Lemma 3.6, we conclude there are embeddings $\psi_G : [m] \rightarrow V(X), \psi_H : [m] \rightarrow V(Y)$. By the choice of V_i , it is clear that $\psi_H(i)$ is adjacent to u_0 whenever $iu \in E(H)$, etc. Furthermore, since $\sigma^{-1}(u_0)\sigma^{-1}(v_0) \in E(\tilde{X})$, we conclude that setting $\psi_H(u) = u_0, \psi_H(v) = v_0, \psi_G(u) = \sigma^{-1}(u_0), \psi_G(v) = \sigma^{-1}(v_0)$ extends ψ_G, ψ_H to the embeddings $\psi_G : V(G^{(\rho)}) \rightarrow V(\tilde{X})$ and $\psi_H : V(H^{(\rho)}) \rightarrow V(Y)$, thus completing the proof. \square

3.6. Completing the proof. Having presented all prerequisites in the previous sections, we now put everything together and present the whole proof of Proposition 3.1, which implies Theorem 1.3.

Proof of Proposition 3.1. Since, with high probability, any arrangement σ can be transformed into a balanced arrangement without ever swapping u_0, v_0 , Proposition 2.2 shows that it suffices to check the statement of Proposition 3.1 only for balanced arrangements σ . If σ satisfies the assumptions of the case (C ρ), for some $\rho \in \{1, 2, 3, 4\}$, combining Proposition 3.8 and Lemma 3.6, one concludes that, with high probability, there exist embeddings $\psi_G : V(G^{(\rho)}) \rightarrow V(X), \psi_H : V(H^{(\rho)}) \rightarrow V(Y)$ with the properties required by Proposition 3.2. Since Lemmas 3.3 and 3.4 show that the vertices $u, v \in V(H^{(\rho)})$ are $(G^{(\rho)}, H^{(\rho)})$ -exchangeable, Proposition 3.2 applies to show that u_0 and v_0 are indeed (\tilde{X}, Y) -exchangeable, which completes the discussion when σ satisfies the conditions of one of the cases (C ρ).

If σ does not satisfy any of these cases, a symmetry argument reduces to the previously discussed case. There are three essential symmetries we use. Firstly, swapping the names of u_0 and v_0 does not affect the fact that these vertices are (X, Y) -exchangeable from σ . This symmetry has the effect of changing A_X and A_Y into B_X and B_Y , and vice versa. Further, noting that u_0, v_0 are exchangeable from σ if and only if they are exchangeable from $(u_0 v_0) \circ \sigma$ allows us to change A_Y into B_Y and vice versa without altering A_X and B_X . Finally, noting that u_0, v_0 are (X, Y) -exchangeable from σ if and only if $\sigma^{-1}(u_0), \sigma^{-1}(v_0)$ are (Y, X) -exchangeable from σ^{-1} allows us to swap labels A_X and B_X into A_Y and B_Y , and vice versa. We apply these symmetries as follows.

If σ majority-maps $N_X(u'_0), N_X(v'_0)$ into the same partite set of Y and σ^{-1} majority-maps $N_Y(u_0), N_Y(v_0)$ into the same partite set of X , we first apply the first symmetry, if needed, in order to have both $N_Y(u_0)$ and $N_Y(v_0)$ majority-mapped into A_X . Then, we apply the second symmetry, if needed, to have both A_X and B_X majority-mapped into B_Y , without affecting the labels of A_X and B_X . This reduces σ to the case (C1).

If σ majority-maps $N_X(u'_0), N_X(v'_0)$ to different partite sets of Y , and σ^{-1} majority-maps $N_Y(u_0), N_Y(v_0)$ to different partite sets of X , we perform a second symmetry, if needed, in order to ensure that $N_X(u'_0)$ is majority-mapped into B_Y and $N_X(v'_0)$ is majority-mapped into A_Y . Depending on where $N_Y(u_0)$ and $N_Y(v_0)$ are mapped, we either reduce to the case (C3) or case (C4).

Finally, if one pair of sets $N_X(u'_0), N_X(v'_0)$ and $N_Y(u_0), N_Y(v_0)$ is majority-mapped into the same partite set while the other pair is majority-mapped into different sets, we apply the third symmetry to ensure that $N_Y(u_0)$ and $N_Y(v_0)$ are mapped to different sets. Then, by applying a combination of the first two symmetries, we can reduce to the case (C2).

The above discussion shows that all cases can be reduced to the four distinguished cases, and hence it suffices to consider only these cases. Since the discussion for these was already presented, the proof is complete. \square

4. GENERALIZATIONS OF WILSON'S THEOREM

The goal of this section is to characterize multiplicity lists $c \in \mathbb{Z}_{>0}^m$ for which the graphs $\text{FS}_m(X, \text{Star}_m)$ and $\text{FS}_m(\text{Star}_n, X)$ are connected, by showing Theorems 1.8 and 1.9. Let us begin by addressing the case in which X is a multiplicity graph.

Proof of Theorem 1.8. Suppose first X is Wilsonian. It is simple to check that the lift X' of a Wilsonian graph X with any multiplicity list $c \in \mathbb{Z}_{>0}^m$ is still Wilsonian, and hence Wilson's theorem implies that the graph $\text{FS}(\text{Star}_n, X')$ is connected. Using Corollary 2.9, this shows that $\text{FS}_m(\text{Star}_n, X)$ is also connected. In the case X is biconnected but not Wilsonian and at least one vertex $v \in V(X)$ has multiplicity $c_v \geq 2$, we will again show that the lift of X is Wilsonian. If v is not an isolated vertex, replacing it with a clique of size at least 2 induces a triangle in the lift graph X' . Since X' is biconnected, and neither cycles of length at least 4 nor the graph θ_0 contain triangles, we conclude that X' must be Wilsonian. As in the previous case, this suffices to show $\text{FS}_m(\text{Star}_n, X)$ is connected. On the other hand, if all multiplicities c_v are 1, the graph $\text{FS}_m(\text{Star}_n, X)$ is isomorphic to the classical version

$\text{FS}(\text{Star}_n, X)$, which has no multiplicities. Wilson's theorem then shows that $\text{FS}(\text{Star}_n, X)$ is not connected, which completes this case.

Finally, we address the case when X is not biconnected. Suppose that all cut vertices of X have multiplicities at least 2. Again, the goal is to show X' is Wilsonian. Since the lift X' contains a triangle, it cannot be either a cycle of length at least 4 or θ_0 . Similarly, since all cut vertices of X have multiplicities at least 2, the graph X' has no cut vertices and hence it is biconnected. Therefore, X' is Wilsonian and $\text{FS}_m(\text{Star}_n, X)$ is biconnected. On the other hand, when X has a cut vertex of multiplicity one, Proposition 2.3 applies because the center of Star_n is a cut vertex. Therefore, we conclude that $\text{FS}_m(\text{Star}_n, X)$ is disconnected, which completes the proof of Theorem 1.8. \square

Now, we focus on the case when Star_n is the multiplicity graph, and prove Theorem 1.9. The proof will be split into three propositions — Proposition 4.1 will consider the case when X contains a k -bridge, Proposition 4.4 will consider the case when X does not contain a k -bridge, and Proposition 4.5 will consider the case when X is a cycle.

Throughout the whole discussion, we will consider vertices $\sigma \in V(\text{FS}_m(X, \text{Star}_m))$ as assignments of labels, represented by vertices of Star_m , to the vertices of X , with the condition that the label $i \in V(\text{Star}_m)$ appears c_i times. Moreover, following Wilson's notation from [13], we denote the center of the star by \emptyset and we denote the multiplicity of \emptyset by $k = c_\emptyset$.

Proposition 4.1. *Let X be a simple graph containing a k -bridge and let Star_m be the star graph, along with the multiplicity list $c \in \mathbb{Z}_{>0}^m$, in which the center has multiplicity $c_\emptyset = k$. Then, the friends-and-strangers graph $\text{FS}_m(X, \text{Star}_m)$ is not connected.*

Proof. Since X contains a k -bridge a_1, \dots, a_k , its subgraph $X|_{V(X) - \{a_2, \dots, a_{k-1}\}}$ splits into connected components A_0 and B_0 , where $a_1 \in A_0$ and $a_k \in B_0$. We also set $A = A_0 - \{a_1\}$, $B = B_0 - \{a_k\}$. As k is fixed throughout the proof, we call this specific k -bridge just the bridge. Without loss of generality, we may assume $|A| \geq |B|$.

Since the star Star_m has at least two leaves, there exists a leaf whose multiplicity is at most $|A|$; let us denote this leaf by l . Consider the assignment $\sigma : V(X) \rightarrow V(\text{Star}_m)$, $\sigma \in V(\text{FS}_m(X, \text{Star}_m))$, which satisfies $\sigma^{-1}(l) \subset A$ and $\sigma^{-1}(\emptyset) = \{a_1, \dots, a_k\}$, and let \mathcal{C} denote the connected component containing σ in $\text{FS}_m(X, \text{Star}_m)$. We will show that any vertex $\tau \in \mathcal{C}$ satisfies $\tau^{-1}(l) \subset A \cup \{a_1, \dots, a_k\}$. Showing this would immediately imply that $\text{FS}_m(X, \text{Star}_m)$ is disconnected, since there are assignments $\tau' \in V(\text{FS}_m(X, \text{Star}_m))$ violating this property.

In fact, we show a stronger statement. For an assignment of labels $\tau : V(X) \rightarrow V(\text{Star}_m)$, we let $x(\tau) = |\tau^{-1}(\emptyset) \cap A|$ and we let $1 \leq i_1 < i_2 < \dots < i_r \leq k$ be the complete set of indices for which $\tau(a_{i_j}) \neq \emptyset$. Letting $I(\tau) = \{i_1, \dots, i_{x(\tau)}\}$, we will show that for any $\tau \in \mathcal{C}$ we have

$$(6) \quad \tau^{-1}(l) \subseteq A \cup \{a_i : i \in I(\tau)\}.$$

Intuitively, $x(\tau)$ is the number of blank symbols assigned to the vertices of A , and we claim that all vertices labeled l under τ appear either in A or among the first $x(\tau)$ non-blank labels on the bridge. Note that there are indeed at least $x(\tau)$ non-blank labels on the bridge, i.e., $r \geq x(\tau)$, because there are at most $k - x(\tau)$ blank labels on the bridge.

It suffices to check that, if τ_1 and τ_2 differ by a (X, Star_m) -friendly swap, then τ_1 satisfies the property (6) if and only if τ_2 also does. Let $I(\tau_1) = \{i_1, \dots, i_{x(\tau_1)}\}$ and $I(\tau_2) = \{j_1, \dots, j_{x(\tau_2)}\}$. Let $uv \in E(X)$ be the edge on which the (X, Star_m) -friendly swap is performed. We consider several cases, based on the position of the edge uv .

Case 1: Suppose that $uv = a_\nu a_{\nu+1}$ for some $\nu \in \{1, \dots, n-1\}$. Since $u, v \notin A$, we have $x(\tau_1) = x(\tau_2)$. As τ_1 and τ_2 are symmetric and one of the swapped labels is \emptyset , we may assume that $\tau_1(a_\nu) = \emptyset = \tau_2(a_{\nu+1})$. Since $\tau_1(a_{\nu+1}) \neq \emptyset$, we must have $\nu+1 = i_t$ for some $t \in \{1, \dots, r\}$. Furthermore, since $\tau_2(a_i) \neq \emptyset$ and τ_1, τ_2 agree on $V(X) - \{a_\nu, a_{\nu+1}\}$, we must also have $\nu = j_t$, and $i_\lambda = j_\lambda$ for $\lambda \neq t$. Since $\tau_1 = \tau_2 \circ (a_i a_{i+1}) = \tau_2 \circ (a_{i_t} a_{j_t})$, it is not hard to see that $\tau_1^{-1}(l) \subseteq A \cup \{a_{i_1}, \dots, a_{i_{x(\tau_1)}}\}$ is equivalent to $\tau_2^{-1}(l) \subseteq A \cup \{a_{j_1}, \dots, a_{j_{x(\tau_2)}}\}$, which is what we aimed to show.

Case 2: Suppose that $u = a_1, v \in A$. If we assume, by symmetry, that $\tau_1(a_1) = \emptyset$, we see that $\tau_2^{-1}(\emptyset) \cap A = (\tau_1^{-1}(\emptyset) \cap A) \cup \{v\}$, implying $x(\tau_2) = x(\tau_1) + 1$. Since $\tau_2(a_1) \neq \emptyset$, we also have $j_1 = 1$ and $j_t = i_{t+1}$ for $t \in \{1, \dots, x(\tau_1)\}$. Since the transposition $(a_1 v)$ fixes the set $A \cup \{a_1, a_{i_1}, \dots, a_{i_{x(\tau_1)}}\}$ and we have $\tau_1^{-1}(a_1) \neq l$, we have the equivalence $\tau_1^{-1}(l) \subseteq A \cup \{a_{i_1}, \dots, a_{i_{x(\tau_1)}}\} \iff \tau_2^{-1}(l) \subseteq A \cup \{a_1, a_{i_1}, \dots, a_{i_{x(\tau_2)}}\}$.

Case 3: Suppose that $u, v \in A$ or $u, v \in B$. Then, $x(\tau_1) = x(\tau_2)$ and $\{i_1, \dots, i_{x(\tau_1)}\} = \{j_1, \dots, j_{x(\tau_2)}\}$. Since the transposition $(u v)$ fixes the set $A \cup \{a_{i_1}, \dots, a_{i_{x(\tau_1)}}\}$, the equivalence is clear.

Case 4: Suppose that $u = a_k, v \in B$. By symmetry, we assume that $\tau_1(a_k) = \emptyset$, and consequently $i_{x(\tau_1)} \leq k-1$, i.e., the last non-blank label is at position at most $k-1$. Since neither of u, v is in A , the number of blank labels in A remains constant and so $x(\tau_2) = x(\tau_1)$. Furthermore, as none of the first $x(\tau_1)$ non-blank labels on the bridge is involved in the swap, we have $\{i_1, \dots, i_{x(\tau_1)}\} = \{j_1, \dots, j_{x(\tau_2)}\}$. Therefore, we also see that the last non-blank

label in τ_2 is at position at most $j_{x(\tau_2)} \leq k - 1$. This means that the transposition $(a_k v)$ involves no elements of $A \cup \{a_{i_1}, \dots, a_{i_x(\tau_1)}\}$, which concludes this case.

Having completed all the cases, we conclude that $\tau^{-1}(l) \subseteq A \cup \{a_{i_1}, \dots, a_{i_x(\tau)}\}$ for all $\tau \in \mathcal{C}$, which shows that the graph $\text{FS}_m(X, \text{Star}_m)$ is not connected when X contains a k -bridge. \square

Remark 4.2. The proof of Proposition 4.1 is reminiscent of Theorem 6.1 from [6], although the technical details of the proof are slightly more general in our proof, due to the multiplicities. Proposition 4.1 can be generalized even further, with almost the same proof, to give the following statement: suppose X is a simple graph containing a k -bridge and Y is a multiplicity graph containing vertices v_1, \dots, v_t whose removal disconnects Y . If $c_{v_1} + \dots + c_{v_t} \leq k$, then $\text{FS}_m(X, Y)$ is disconnected.

Now, we focus on showing that the graph $\text{FS}_m(X, \text{Star}_m)$ is connected whenever X has no k -bridges, which is shown in Proposition 4.4. Before presenting the proof of this proposition, we present two simple results which will be used throughout the proof.

Lemma 4.3. *Suppose that X is a connected simple graph on $k + 2$ vertices, not isomorphic to a path or a cycle. Then, the graph $\text{FS}_m(X, \text{Star}_m)$ is connected.*

Proof. By Proposition 2.1, it suffices to show that any two adjacent vertices $u, v \in V(X)$ are (X, Star_m) -exchangeable from any assignment $\sigma \in V(\text{FS}_m(X, \text{Star}_m))$. This is clear if any of these vertices is labeled by \emptyset , or if $\sigma(u) = \sigma(v)$. Hence, let us assume that $\emptyset \neq \sigma(u) \neq \sigma(v) \neq \emptyset$. Note that all vertices except u and v are labeled by \emptyset under σ . Since X is not a path or a cycle, it contains a vertex of degree at least 3 — let x_1 be this vertex. By symmetry of u and v , we may assume that there is a path from v to x_1 with vertices $v, u, x_t, \dots, x_2, x_1$. Suppose also that x_1 has neighbors y and z , which are different from x_2 . Then, the sequence of swaps $ux_t, \dots, x_2x_1, x_1y, vu, ux_t, \dots, x_2x_1, x_1z, x_1y, x_1x_2, \dots, x_tu, uv, zx_1, x_1x_2, \dots, x_tu$ exchanges u and v . \square

Proposition 4.4. *Suppose X is a simple graph that does not contain a k -bridge and is not isomorphic to a cycle. Then, the graph $\text{FS}_m(X, \text{Star}_m)$ is connected.*

Proof. Our main tool will be the notion of exchangeability, which was defined in Section 2. Namely, Proposition 2.1 reduces showing that $\text{FS}_m(X, \text{Star}_m)$ is connected to showing that any two adjacent vertices $u, v \in V(X)$ are (X, Star_m) -exchangeable from any starting arrangement $\sigma \in V(\text{FS}_m(X, \text{Star}_m))$, which is done by explicitly constructing a sequence of (X, Star_m) -friendly swaps which transforms σ into $\sigma \circ (uv)$. If either $\sigma(u) = \emptyset$ or $\sigma(v) = \emptyset$, it is easy to see that the swap along the edge uv itself is a (X, Star_m) -friendly swap, thus accomplishing the goal. Therefore, may we assume that $\sigma(u), \sigma(v) \neq \emptyset$ in the rest of the proof.

Following the notation from [13], for a path $p = a_1a_2 \cdots a_l$ and an assignment τ with $\tau^{-1}(a_1) = \emptyset$, we denote the sequence of swaps along the edges $a_1a_2, a_2a_3, \dots, a_{l-1}a_l$ by σ_p . Note that τ and $\tau \circ \sigma_p$ are in the same connected component of $\text{FS}(X, \text{Star}_m)$.

The first step of the proof will be *move* every blank symbol from its original position to obtain an arrangement in which the set of vertices labeled by a blank symbol, together with u and v , forms a connected subgraph of X . More formally, we will show there exists a sequence of (X, Star_m) -friendly swaps, not including vertices u and v , which transforms σ into an assignment σ' for which $X|_{\sigma'^{-1}(\emptyset) \cup \{u, v\}}$ is connected. After ensuring this subgraph is not a path or a cycle, we will be in the position to apply Lemma 4.3, which will show that u and v are exchangeable.

We will achieve our first goal inductively. Suppose that after several swaps, the vertices u, v , and i other vertices with blank labels form a connected subgraph $C \subset X$. If $i < k$, pick a vertex $w \in V(X)$ labeled by \emptyset and let p be the shortest path from w to a vertex in the neighborhood of C . Applying the sequence of swaps σ_p transforms the existing arrangement into an arrangement in which the vertices u, v , and $i + 1$ other vertices with blank labels form a connected subgraph of X . If we apply this procedure as long as $i < k$, we arrive at the arrangement σ' with the property that $X|_{\sigma'^{-1}(\emptyset) \cup \{u, v\}}$ is connected.

Hence, $X|_{\sigma'^{-1}(\emptyset) \cup \{u, v\}}$ is a connected subgraph of X on $k + 2$ vertices, and we denote it by H . Since vertices u, v were not included in any of the swaps transforming σ into σ_k , by Proposition 2.2, it suffices to show that the pair u, v is (X, Star_m) -exchangeable from σ' .

In the case H is not isomorphic to a path, Lemma 4.3 applies to subgraphs of X and Star_m spanned by the vertices $\{u, v\} \cup \sigma'^{-1}(\emptyset)$ and $\{\sigma(u), \sigma(v), \emptyset\}$, and it shows that the vertices u and v are exchangeable with respect to these subgraphs. However, this also means u, v are (X, Star_m) -exchangeable from σ' . Since transforming σ into σ' did not involve swaps including u or v , Proposition 2.2 shows u, v are (X, Star_m) -exchangeable from σ .

In the case H is isomorphic to a path or a cycle, we denote its vertices by a_1, \dots, a_{k+2} , having $u = a_i$ and $v = a_{i+1}$. We now consider two cases, based on whether any vertex of H has degree 3 in X .

Case 1: Suppose that a vertex of $a_j \in H$ has degree at least 3 in X . Then, we either have $a_j \notin \{u, v\}$ or $a_j \in \{u, v\}$. In the case $a_j \notin \{u, v\}$, we may assume that $j < i$, by reversing the indices of a_1, \dots, a_k . Then, the sequence of swaps $ya_j, a_j a_{j-1}, \dots, a_2 a_1$ transforms σ_k into σ'_k , in which $X|_{\sigma'_k^{-1}(\emptyset) \cup \{u\}}$ is not isomorphic to a path, thus reducing this case to Lemma 4.3.

Suppose now that $a_j \in \{u, v\}$. By potentially relabeling u and v or reversing indices of a_1, \dots, a_k , we may assume that $a_j = a_i = u$ and $a_{i+1} = v$. Since $i = j \leq k - 2$, we have $i + 2 \leq k$ and we may consider the following sequence of swaps

$$(7) \quad a_{i+2}v, vu, a_i y, a_{i-1}u, a_{i-1}a_{i-2}, \dots, a_2 a_1.$$

This sequence of swaps transforms σ_k into σ'_k in which $X|_{\sigma'_k^{-1}(\emptyset) \cup \{u\}}$ is not isomorphic to a path, thus reducing this case to Lemma 4.3. After exchanging u, v using Lemma 4.3, it suffices to reverse the sequence (7) to show u, v are (X, Star_m) -exchangeable from σ' .

Case 2: Suppose that H does not contain a vertex of degree 3 in X . Since X is connected, H is not a cycle, and because X does not contain a k -bridge, the graph $X - \{a_3, \dots, a_k\}$ is connected. Hence, H must be an interval on the cycle C , whose vertices we denote by $a_1, \dots, a_{k+2}, a_{k+3}, \dots, a_l$. Since X is not a cycle, there exists a vertex $a_t \in C$ of degree at least 3 in X . If we pick any vertex a_j of H labeled by \emptyset , performing swaps along edges $a_j a_{j+1}, \dots, a_{l-1} a_l, a_1 a_1, \dots, a_{j-2} a_{j-1}$ has the effect of shifting all labels of C one position backwards. By performing this sequence of swaps a sufficient number of times, we may assume that the labels $\sigma(u)$ and $\sigma(v)$ are moved to vertices a_t and a_{t+1} . Now, a vertex labeled by $\sigma(u)$ has degree at least 3, and hence Case 1 implies that a_t and a_{t+1} are (X, Star_m) -exchangeable from the resulting arrangement. Reversing the counterclockwise shifts previously performed shows that $\sigma' \circ (u v)$ can be obtained from σ' through a sequence of (X, Star_m) -friendly swaps, which completes the proof. \square

Finally, to complete the discussion of $\text{FS}_m(X, \text{Star}_m)$, we address the case when X is a cycle.

Proposition 4.5. *Let Star_m be a multiplicity graph with the multiplicity list $c \in \mathbb{Z}_{>0}^m$ with $c_\emptyset = k$ and total multiplicity n . Then, $\text{FS}_m(\text{Cycle}_n, \text{Star}_m)$ is connected if and only if $m = 3$ and one of the leaves of Star_m has multiplicity 1.*

Proof. The key observation in this proof is that the cyclic order of non-blank labels remains unchanged after a $(\text{Cycle}_n, \text{Star}_m)$ -friendly swap. Hence, if $\text{FS}(\text{Cycle}_n, \text{Star}_m)$ is connected, there must be only one possible cyclic order of non-blank labels. Suppose that the multiplicities of the leaves of Star_m are c_1, \dots, c_{m-1} . The number of cyclic orders of elements $1, \dots, m-1$ where element i appears with multiplicity c_i is $\frac{(c_1 + \dots + c_{m-1} - 1)!}{c_1! \dots c_{m-1}!}$, which equals 1 if and only if $m = 3$ and either $c_1 = 1$ or $c_2 = 1$. On the other hand, if the cyclic order of non-blank labels matches in two different assignments, it is not hard to see that one can be transformed into the other using a sequence of $(\text{Cycle}_n, \text{Star}_m)$ -friendly swaps. \square

The combination of Propositions 4.1, 4.4 and 4.5 immediately yields the proof of Theorem 1.9, completing this section.

5. PATHS AND CYCLES

5.1. Background and terminology. In this section, we investigate the structure of the friends-and-strangers graphs arising when one of the graphs is a path or a cycle, while the other graph is a multiplicity graph. This extends the results obtained by Defant and Kravitz in [6], who characterized the connected components of $\text{FS}(\text{Path}_n, Y)$, $\text{FS}(\text{Cycle}_n, Y)$ for a simple graph Y in terms of the acyclic orientations of its complement \bar{Y} . Before presenting the generalizations, we introduce relevant terminology and make a brief overview of the results obtained by Defant and Kravitz.

For a simple graph G , an *orientation* of the graph corresponds to choosing a direction for each of its edges, and an orientation is *acyclic* if the arising directed graph has no cycles. We denote the set of acyclic orientations of G by $\text{Acyc}(G)$. Further, if we think of Path_n and Cycle_n as graphs on the vertex set $[n]$, we see that every bijection $\sigma : [n] \rightarrow V(G)$ induces an acyclic orientation $\alpha_G(\sigma)$ in which an edge $uv \in E(G)$ is directed from $u \rightarrow v$ if and only if $\sigma^{-1}(u) < \sigma^{-1}(v)$. Finally, for a given acyclic orientation $\alpha \in \text{Acyc}(G)$, we define the set of its *linear extensions* $\mathcal{L}(\alpha)$ to be the set of all bijections $\sigma : [n] \rightarrow V(G)$ which induce the orientation α on G , i.e., $\mathcal{L}(\alpha) = \{\sigma : [n] \rightarrow V(G) \mid \alpha_G(\sigma) = \alpha\}$. Even more generally, for a set of acyclic orientations $A \subseteq \text{Acyc}(G)$, we may define the set of linear extensions of A to be $\mathcal{L}(A) = \bigcup_{\alpha \in A} \mathcal{L}(\alpha)$. In [6], Defant and Kravitz showed that arrangements $\sigma, \tau : V(\text{Path}_n) \rightarrow V(Y)$ are in the same connected component of $\text{FS}(\text{Path}_n, Y)$ if and only if they induce the same acyclic orientation of the complement of Y , i.e., $\alpha_{\bar{Y}}(\sigma) = \alpha_{\bar{Y}}(\tau)$.

More work is needed to describe the connected components of Cycle_n , since we need to introduce the notions of toric and double flip equivalence. For a given acyclic $\alpha \in \text{Acyc}(G)$, a vertex $u \in V(G)$ is a *source* of α if all edges $uv \in E(G)$ are oriented $u \rightarrow v$, and a vertex $u \in V(G)$ is a *sink* of α if all edges $uv \in E(G)$ are oriented $v \rightarrow u$. Given an orientation $\alpha \in \text{Acyc}(G)$ and a vertex u which is either a source or a sink of α , reversing all edges incident to u gives a new acyclic orientation of G . We refer to this operation as a *flip*, and which is called *positive* if u is a source of α , and *negative* otherwise. This allows to define an equivalence relation called *toric equivalence* by declaring that α, α' are equivalent if one can be obtained from the other by a sequence of flips. If α, α' are torically equivalent, we write $\alpha \sim \alpha'$. Also, we write $\text{Acyc}(G)/\sim$ for the set of toric equivalence classes of acyclic orientations, and we call elements of $\text{Acyc}(G)/\sim$ *toric acyclic orientations*.

If u, v are non-adjacent vertices of G such that u is a source and v is a sink of α , respectively, performing two simultaneous flips on α , at u and v , produces a new acyclic orientation. We call this operation a *double flip*, and, as before, we define the equivalence relation \approx , called *double flip equivalence*, by declaring that two acyclic orientations are double flip equivalent if one can be obtained from the other by a sequence of double flips. It is worth noting that double flip equivalence implies toric equivalence of two orientations, whereas the reverse is not always true. We denote the set of double flip equivalence classes by $\text{Acyc}(G)/\approx$ and we denote the equivalence class of an orientation $\alpha \in \text{Acyc}(G)$ by $[\alpha]_{\approx}$.

In this language, one can describe the connected components of $\text{FS}(\text{Cycle}_n, Y)$ in several ways. Theorem 4.1 of [6] states that the arrangements $\sigma, \tau : V(\text{Cycle}_n) \rightarrow V(Y)$ are in the same connected component of $\text{FS}(\text{Cycle}_n, Y)$ if and only if their induced acyclic orientations of \overline{Y} are double flip equivalent, i.e., $\alpha_{\overline{Y}}(\sigma) \approx \alpha_{\overline{Y}}(\tau)$. Investigating the relation between double flip equivalence and toric equivalence further, Defant and Kravitz showed that every toric equivalence class $[\alpha]_{\sim}$ can be viewed as a union of double flip equivalence classes forming an orbit under an action of a certain abelian group. This allows them to characterize all graphs Y for which $\text{FS}(\text{Cycle}_n, Y)$ is connected, and the answer turns out to be the set of all graphs Y whose complement is a forest of trees of coprime sizes.

In the process of showing these results, Defant and Kravitz also showed the following useful fact, which we will also make use of. Namely, Corollary 4.2 of [6] shows that if $\alpha \approx \beta$ and $\vec{\alpha}, \vec{\beta}$ arise from α, β , respectively, by flipping one source into a sink, then $\vec{\alpha} \approx \vec{\beta}$. The following statement is a direct corollary of their result.

Corollary 5.1. *Let $\alpha, \beta \in \text{Acyc}(G)$ be acyclic orientations, and assume that β is obtained from α through a sequence of a_+ positive flips and a_- negative flips. Then, any sequence of $a_+ - a_-$ positive flips transforms α into β' , which is double flip equivalent to β .*

In the rest of this section, we present a generalization of results from [6] to the multiplicity setting. First, we discuss the graphs $\text{FS}_m(\text{Path}_n, X)$, and show that their connected components correspond to equivalence classes of $\text{Acyc}(\overline{X'})$ under a certain equivalence relation. Then, we characterize the connected components of $\text{FS}_m(\text{Cycle}_n, X)$ and give a full characterization of graphs X for which $\text{FS}_m(\text{Cycle}_n, X)$ is connected.

5.2. Connectivity of $\text{FS}_m(\text{Path}_n, X)$. We begin by describing the connected components of $\text{FS}_m(\text{Path}_n, X)$ for a multiplicity graph X , which is analogous to Theorem 3.1 of [6]. The main technical tool we use is the notion of permutation equivalence and the bulk of work lies in investigating how this equivalence relation interacts with toric and double flip equivalence on the lift of X .

More precisely, the notion of permutation equivalence \equiv , introduced in Section 2, can be extended to acyclic orientations in the following way. If X' is the lift of X , having the clique decomposition $V(X') = \bigsqcup_{v \in V(X)} S_v$, we can define the action of \mathfrak{S}_X on $\text{Acyc}(X')$ by setting $\rho(\alpha_{\overline{X'}}(\sigma)) = \alpha_{\overline{X'}}(\rho \circ \sigma)$ for every $\sigma : [n] \rightarrow V(X')$. Recall that $\mathfrak{S}_X = \prod_{v \in V(X)} \mathfrak{S}_{S_v}$ denotes the set of permutation which fix the cliques S_v of X' .

We say that acyclic orientations $\alpha, \beta \in \text{Acyc}(\overline{X'})$ are *permutation equivalent*, written as $\alpha \equiv \beta$, if they are in the same orbit of the action of \mathfrak{S}_X on $\text{Acyc}(X')$. In combinatorial terms, we have $\alpha \equiv \beta$ precisely when there is a permutation $\rho \in \mathfrak{S}_X$ which transforms α into β , in the sense that $u \rightarrow_{\alpha} v$ is equivalent to $\rho(u) \rightarrow_{\beta} \rho(v)$. A simple consequence of this definition is that permutation equivalent arrangements $\sigma, \tau : V(\text{Cycle}_n) \rightarrow V(X')$ induce permutation equivalent orientations $\alpha_{\overline{X'}}(\sigma) \equiv \alpha_{\overline{X'}}(\tau)$.

Theorem 5.2. *Let X be a multiplicity graph of total multiplicity n and let X' be its lift. The arrangements $[\sigma]_{\equiv}, [\tau]_{\equiv} \in V(\text{FS}_m(\text{Path}_n, X))$ are in the same connected component if and only if σ, τ induce permutation equivalent orientations on $\overline{X'}$, i.e. $\alpha_{\overline{X'}}(\sigma) \equiv \alpha_{\overline{X'}}(\tau)$.*

Proof. In this proof, we use Corollary 2.10 to reduce Theorem 5.2 to the Theorem 3.1 of [6].

From Corollary 2.10, we know that $[\sigma]_{\equiv}, [\tau]_{\equiv}$ are in the same connected component of $\text{FS}_m(\text{Path}_n, X)$ if and only if σ, τ are in the same connected component of G , where G is the graph $\text{FS}(\text{Path}_n, X')$ with the added set of edges between pairs of permutation equivalent arrangements.

Suppose now that σ, τ are indeed in the same connected component of G . To show $\alpha_{\overline{X\tau}}(\sigma) \equiv \alpha_{\overline{X\tau}}(\tau)$, it suffices to check this in case σ and τ are adjacent. In case σ, τ differ by a (Path_n, X') -friendly swap, it follows from Theorem 3.1 of [6] that $\alpha_{\overline{X\tau}}(\sigma) = \alpha_{\overline{X\tau}}(\tau)$. In case σ and τ are permutation equivalent, it follows from our previous observation that so their induced orientations must be permutation equivalent as well.

To show the other direction, assume $\alpha_{\overline{X\tau}}(\tau) \equiv \alpha_{\overline{X\tau}}(\sigma)$. Then there exists a permutation $\rho \in \mathfrak{S}_X$ such that $\alpha_{\overline{X\tau}}(\tau) = \alpha_{\overline{X\tau}}(\rho \circ \sigma)$. Theorem 3.1 of [6] ensures that τ and $\rho \circ \sigma$ are connected in $\text{FS}(\text{Path}_n, X')$, while $\sigma \equiv \rho \circ \sigma$, implying $\{\sigma, \rho \circ \sigma\} \in E_0$. Combining these, we see σ is connected to τ in G , completing the proof. \square

5.3. Connectivity of $\text{FS}_m(\text{Cycle}_n, X)$. We now focus on describing the connected components of $\text{FS}_m(\text{Cycle}_n, X)$. The proofs in this section add two new ingredients to the work of Defant and Kravitz. The first new idea is to combine pairs of equivalence relations \sim, \equiv and \approx, \cong into their finest common coarsening, while the second one is to define the period of an acyclic orientation $\alpha \in \text{Acyc}(\overline{X'})$.

More precisely, the first step is to define an equivalence relation \simeq on $\text{Acyc}(\overline{X'})$ in which $\alpha \simeq \beta$ if and only if there exists a sequence of orientations $\alpha = \alpha_0, \alpha_1, \dots, \alpha_n = \beta$ in which α_i and α_{i+1} are either permutation equivalent or torically equivalent. Note that permuting moves and flips commute, in the sense that flipping a vertex u in α and obtaining an orientation $\beta \in \text{Acyc}(\overline{X'})$ and then taking β to $\rho(\beta)$ corresponds to flipping a vertex $\rho(u)$ in $\rho(\alpha)$, which also results in the orientation $\rho(\beta)$. Hence, another way to define \simeq is to declare $\alpha \simeq \beta$ whenever there exists an orientation $\gamma \in \text{Acyc}(\overline{X'})$ satisfying $\alpha \equiv \gamma \sim \beta$.

Following the same strategy, one can also define an equivalence relation \cong as the finest common coarsening of \approx and \equiv . Since \simeq and \cong are both coarser than \equiv , they extend to equivalence relations on $\text{Acyc}(\overline{X'})/\equiv$. Using these notions, we can generalize Theorem 4.1 of [6] in the following way.

Theorem 5.3. *Let X be a multiplicity graph and let X' be its lift. The arrangements $[\sigma]_{\equiv}, [\tau]_{\equiv} \in V(\text{FS}_m(\text{Cycle}_n, X))$ are in the same connected component if and only if $\alpha_{\overline{X\tau}}(\sigma) \cong \alpha_{\overline{X\tau}}(\tau)$.*

Proof. The structure of the proof closely follows the proof of the Theorem 5.2. Recall that Corollary 2.10 guarantees that the connected components of $\text{FS}(X, \text{Cycle}_n)$ correspond to projections under $\sigma \mapsto [\sigma]_{\equiv}$ of connected components in $G = (V(\text{FS}(X', \text{Cycle}_n)), E(\text{FS}(X', \text{Cycle}_n)) \cup E_0)$, where $E_0 = \{\sigma\tau : \sigma, \tau \in V(\text{FS}(X', \text{Cycle}_n)), \sigma \equiv \tau\}$. Hence, to show Theorem 5.3, it suffices to check that $\sigma, \tau \in V(G)$ are connected in G if and only if their induced orientations $\alpha_{\overline{X\tau}}(\sigma)$ and $\alpha_{\overline{X\tau}}(\tau)$ satisfy $\alpha_{\overline{X\tau}}(\sigma) \cong \alpha_{\overline{X\tau}}(\tau)$.

To show the forward direction, assume that σ and τ are in the same connected component. Furthermore, if we induct on the distance between σ and τ , without loss of generality we may assume $\sigma\tau \in E(G) = E(\text{FS}(X', \text{Cycle}_n)) \cup E_0$. If $\sigma\tau \in E_0$, we have $\sigma \equiv \tau$, and hence $\alpha_{\overline{X\tau}}(\sigma) \equiv \alpha_{\overline{X\tau}}(\tau)$. If $\sigma\tau \in E(\text{FS}(X', Y'))$, we have $\alpha_{\overline{X\tau}}(\sigma) \approx \alpha_{\overline{X\tau}}(\tau)$, from Theorem 4.1 of [6]. In either case, we clearly have $\alpha_{\overline{X\tau}}(\sigma) \cong \alpha_{\overline{X\tau}}(\tau)$.

On the other hand, if $\alpha_{\overline{X\tau}}(\sigma) \cong \alpha_{\overline{X\tau}}(\tau)$, then there exists a permutation $\rho \in \mathfrak{S}_X$ and a sequence of double flip moves transforming $\alpha_{\overline{X\tau}}(\sigma)$ into $\alpha_{\overline{X\tau}}(\tau)$. As double flips correspond to edges of $\text{FS}(X', \text{Cycle}_n)$, and permuting with ρ corresponds to an edge in E_0 , we conclude that σ and τ are connected in G , completing the proof. \square

One of the key insights Defant and Kravitz used to understand the structure of $\text{FS}(\text{Cycle}_n, Y)$ was the automorphism $\varphi_n : [n] \rightarrow [n]$ of the graph Cycle_n given by $\varphi_n(k) = k + 1$ for $k < n$ and $\varphi_n(n) = 1$. Intuitively, one can think of φ_n as a cyclically shifting Cycle_n one place forward. Since φ_n is an automorphism of the cycle Cycle_n , it also induces an automorphism φ_n^* of the friends-and-strangers graph $\text{FS}(\text{Cycle}_n, Y)$ (Proposition 2.3 of [6]). It is not hard to see that these ideas extend to the multiplicity setting and that φ_n^* can be viewed both as an automorphism of $\text{FS}(\text{Cycle}_n, X')$ and of $\text{FS}_m(\text{Cycle}_n, X)$, which will be used extensively henceforth. For simplicity of notation, we will suppress the index n when it is clear what the underlying cycle is. Using the automorphism φ^* , we define the second main notion of this section, the period of an orientation.

Definition 5.4. Let $\sigma : [n] \rightarrow V(X')$ be an arrangement corresponding to a vertex of $\text{FS}(\text{Cycle}_n, X')$. The *period* of σ is the smallest positive integer π_σ for which $\sigma \circ \varphi^{\pi_\sigma} \equiv \sigma$. The *period* of an acyclic orientation $\alpha \in \text{Acyc}(\overline{X'})$ is equal to the smallest period π_σ of any linear extension $\sigma \in \mathcal{L}(\alpha)$.

Example 5.5. Consider a multiplicity graph X with $V(X) = \{u_1, u_2, u_3\}$ and $E(X) = \{u_1u_2, u_2u_3\}$. Moreover, assume that u_1 and u_2 have capacity 2 and u_3 has capacity 4. Then, the arrangement $\sigma : [8] \rightarrow V(X)$ sending $1, 5 \mapsto u_1, 3, 7 \mapsto u_2$ and $2, 4, 6, 8 \mapsto u_3$ has period $\pi_\sigma = 4$. On the other hand, the arrangement $\tau : [8] \rightarrow V(X)$ given by $2, 5 \mapsto u_1, 3, 7 \mapsto u_2, 1, 4, 6, 8 \mapsto u_3$ has period 8, even though it induces the same orientation α on $\overline{X'}$ as σ . Moreover, the period of the induced orientation α is 4 in this case.

Proposition 5.6. *Let $\sigma : [n] \rightarrow V(X')$ be an arrangement of period π_σ , and let $V(X') = \bigsqcup_{v \in V(X)} S_v$ be the clique decomposition of $V(X')$. Consider the partition of $[n]$ into congruence classes modulo π_σ , given by $[n] =$*

$\bigsqcup_{j \in [\pi_\sigma]} \{k \in [n] : k \equiv_{\pi_\sigma} j\} = \bigsqcup_{j \in [\pi_\sigma]} P_j$. For every $j \in [\pi_\sigma]$, $\sigma(P_j)$ is fully contained in one of the cliques and each clique S_v contains the images of exactly $c_v n / \pi_\sigma$ congruence classes P_j .

Proof. We begin by showing that, if $\sigma(j) \in S_v$, then $\sigma(j + \pi_\sigma) \in S_v$ as well. From the definition of π_σ , we have $\sigma \circ \varphi^{\pi_\sigma} = \rho \circ \sigma$, for some $\rho \in \mathfrak{S}_X$. Substituting j in this equality gives $\sigma(j + \pi_\sigma) = \rho(\sigma(j))$. However, as ρ only permutes the elements within the cliques, and $\sigma(j) \in S_v$, we infer $\sigma(j + \pi_\sigma) \in S_v$. Repeated application of the above argument now gives that $j \in S_v$ implies $\sigma(P_j) \subseteq S_v$, which completes the proof of the first statement. The second statement follows by comparing cardinalities of sets $\sigma(P_j)$ and S_v . \square

The following proposition shows that the period of torically or permutation equivalent orientations is the same. Hence, we may talk about the period of a equivalence class $[\alpha]_{\simeq} \in \text{Acyc}(\overline{X'}) / \simeq$.

Proposition 5.7. *Let X be a connected graph and let $\alpha, \beta \in \text{Acyc}(\overline{X'})$ be acyclic orientation satisfying $\alpha \simeq \beta$. Then $\pi_\alpha = \pi_\beta$.*

Proof. It suffices to check that $\pi_\alpha = \pi_\beta$ whenever $\alpha \equiv \beta$ and $\alpha \sim \beta$. Let us begin by addressing the case when $\alpha \equiv \beta$. We will show that $\pi_\alpha \geq \pi_\beta$. By symmetry of α and β , this will also imply $\pi_\beta \geq \pi_\alpha$, which suffices to conclude $\pi_\alpha = \pi_\beta$.

As $\alpha \equiv \beta$, we have a permutation $\rho \in \mathfrak{S}_X$ transforming α into β . Let σ be a linear extension of α with period π_α . Then, $\rho \circ \sigma$ is a linear extension of β , with period at most π_α , because $(\rho \circ \sigma) \circ \varphi^{\pi_\alpha} = \rho \circ (\sigma \circ \varphi^{\pi_\alpha}) = \rho \circ \sigma$. Hence, $\pi_\beta \leq \pi_\alpha$, which completes the discussion in the case $\alpha \equiv \beta$.

If $\alpha \sim \beta$, the proof is more involved and needs the following auxiliary lemma.

Lemma 5.8. *Let $\alpha \in \text{Acyc}(\overline{X'})$, and let u be a source of α . Then, there exists a linear extension $\sigma \in \mathcal{L}(\alpha)$ with $\pi_\sigma = \pi_\alpha$ and $\sigma(1) = u$. Similarly, if v is a sink of α , there exists a linear extension $\sigma \in \mathcal{L}(\alpha)$ with $\pi_\sigma = \pi_\alpha$ and $\sigma(n) = v$.*

Proof. As the two statements given in the lemma are completely analogous, we discuss only the case when u is a source. Let $\tau : [n] \rightarrow V(X')$ be a linear extension of α with $\pi_\tau = \pi_\alpha$, which does not necessarily satisfy $\tau(1) = u$.

Consider the set of vertices of X' that are labeled by a number not larger than π_α , which is given by $S = \{\tau(i) : i \leq \pi_\alpha\}$. If we recall the decomposition $[n] = \bigsqcup_{j \in [\pi_\sigma]} P_j$ from Proposition 5.6, we set S to be the set of images of minimal elements of progressions P_j .

We note that $u \in S$. To see why, let S_v be the clique of X' which contains u . As the vertex u is a source, it is the minimum of σ^{-1} on the clique S_v . By Proposition 5.6, we know that the minimum of σ^{-1} on each clique is at most π_σ , as the preimage of every clique contains at least one congruence class modulo π_σ , which implies $u \in S$.

The restriction of α onto S still induces an acyclic orientation, of which $\tau|_{[\pi_\alpha]}$ is still a linear extension. As u is a source in $\alpha|_S$, we infer there exists a linear extension $\sigma : [\pi_\alpha] \rightarrow S$ of $\alpha|_S$ which has $\sigma(1) = u$.

Then, we extend σ to the $\tilde{\sigma} : [n] \rightarrow V(X')$ by setting $\tilde{\sigma}(k\pi_\alpha + j) = \tau(k\pi_\alpha + \tau^{-1} \circ \sigma(j))$, where j ranges through $[\pi_\alpha]$ and $k \in \{0, \dots, n/\pi_\alpha - 1\}$. It is not hard to check that $\tilde{\sigma}$ indeed agrees with σ on $[\pi_\alpha]$, as for $k = 0$ we have $\tilde{\sigma}(j) = \tau \circ \tau^{-1} \circ \sigma(j) = \sigma(j)$. Hence, we have $\tilde{\sigma}(1) = \sigma(1) = u$. It remains to check that $\tilde{\sigma}$ is a proper linear extension of α and that $\tilde{\sigma}$ has period π_α .

Let us first check that $\tilde{\sigma} \in \mathcal{L}(\alpha)$. Suppose $u_1 u_2$ is an edge of $\overline{X'}$, directed from u_1 to u_2 under α , and let $u_1 = \tau(k_1\pi_\alpha + l_1), u_2 = \tau(k_2\pi_\alpha + l_2)$. As $\tau \in \mathcal{L}(\alpha)$, we have

$$(8) \quad k_1\pi_\alpha + l_1 < k_2\pi_\alpha + l_2.$$

The goal is now to show $\tilde{\sigma}^{-1}(u_1) < \tilde{\sigma}^{-1}(u_2)$. From the definition of $\tilde{\sigma}$, we know that $u_1 = \tilde{\sigma}(k_1\pi_\alpha + j_1), u_2 = \tilde{\sigma}(k_2\pi_\alpha + j_2)$, where $j_1 = \sigma^{-1} \circ \tau(l_1), j_2 = \sigma^{-1} \circ \tau(l_2)$. To show that $\tilde{\sigma} \in \mathcal{L}(\alpha)$, we need to verify that

$$(9) \quad k_1\pi_\alpha + j_1 < k_2\pi_\alpha + j_2.$$

From (8), we see that $k_1 \leq k_2$. If $k_1 < k_2$, the inequality (9) follows immediately, as $j_1, j_2 \in \{1, \dots, \pi_\alpha\}$. On the other hand, if $k_1 = k_2$, we have $l_1 \leq l_2$. Hence, there is an edge in $\alpha|_S$ from $\tau(l_1)$ to $\tau(l_2)$. As $\sigma \in \mathcal{L}(\alpha|_S)$, we must also have $\sigma^{-1}(\tau(l_1)) < \sigma^{-1}(\tau(l_2))$, implying $j_1 < j_2$. Thus, $\tilde{\sigma}$ is a linear extension of α .

Now, we show that the period of $\tilde{\sigma}$ is π_α . We have

$$\begin{aligned} \tilde{\sigma} \circ \varphi^{\pi_\alpha}(k\pi_\alpha + j) &= \tilde{\sigma}((k+1)\pi_\alpha + j \pmod n) \\ &= \tau((k+1)\pi_\alpha + \tau^{-1} \circ \sigma(j) \pmod n) \\ &= \tau \circ \varphi^{\pi_\alpha}(k\pi_\alpha + \tau^{-1} \circ \sigma(j)). \end{aligned}$$

As the period of τ is π_α , we infer that $\tau \circ \varphi^{\pi_\alpha}(k\pi_\alpha + \tau^{-1} \circ \sigma(j))$ is in the same clique as $\tau(k\pi_\alpha + \tau^{-1} \circ \sigma(j)) = \tilde{\sigma}(k\pi_\alpha + j)$, meaning that $\tilde{\sigma} \circ \varphi^{\pi_\alpha}(k\pi_\alpha + j)$ and $\tilde{\sigma}(k\pi_\alpha + j)$ are in the same clique. This shows that the period of $\tilde{\sigma}$ is at most π_α , but since $\tilde{\sigma}$ is a linear extension of α , the period of $\tilde{\sigma}$ is exactly π_α . \square

Now, we are ready to show that the period of an orientation is preserved under toric equivalence. If $\alpha \sim \beta$, the orientation β can be obtained from α by applying a sequence of flip moves. By induction, it suffices to address the case when α and β differ by a single flip on the vertex u , where we may assume that u is a sink in α and a source in β . Using Lemma 5.8, we find a linear extension $\sigma \in \mathcal{L}(\alpha)$ for which $\pi_\sigma = \pi_\alpha$ and $\sigma(n) = u$. As α and β differ by a flip of u , we see that $\sigma \circ \varphi$ is a linear extension of β .

Note that $\sigma \circ \varphi$ has period as most π_α , because $\sigma \circ \varphi \circ \varphi^{\pi_\alpha} = \sigma \circ \varphi^{\pi_\alpha} \circ \varphi \equiv \sigma \circ \varphi$. Hence, $\pi_\beta \leq \pi_\alpha$, which suffices to complete the proof by applying symmetry. \square

In Theorem 4.7 of [6], Defant and Kravitz obtain the following characterization of the connected components of $\text{FS}(\text{Cycle}_n, Y)$.

Proposition 5.9 ([6]). *Let Y be a simple graph on the vertex set $[n]$. Let n_1, \dots, n_r denote the sizes of the connected components of \overline{Y} , and let $\nu = \gcd(n_1, \dots, n_r)$. For each toric acyclic orientation $[\alpha]_\sim \in \text{Acyc}(\overline{Y})/\sim$, choose a linear extension $\sigma_{[\alpha]_\sim}$, and let $J_{[\alpha]_\sim}$ be the connected component of $\text{FS}(\text{Cycle}_n, Y)$ containing $[\sigma_{[\alpha]_\sim}]$. Then*

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

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

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

We obtain the following generalization of this result.

Theorem 5.10. *Let X be a multiplicity graph and suppose $\overline{X'}$ has connected components Z_1, \dots, Z_r . An orientation $[\alpha]_{\simeq} \in \text{Acyc}(\overline{X'})/\simeq$ induces orientations $[\alpha^{(1)}]_{\simeq}, \dots, [\alpha^{(r)}]_{\simeq}$ on Z_1, \dots, Z_r . Let π_1, \dots, π_r be the periods of the induced orientations, and let $\delta_\alpha = \gcd(\pi_1, \dots, \pi_r)$. For each toric acyclic orientation $[\alpha]_{\simeq} \in \text{Acyc}(\overline{X'})/\simeq$, choose a linear extension $\sigma_{[\alpha]_{\simeq}} : [n] \rightarrow V(X')$, and let $H_{[\alpha]_{\simeq}}$ be the connected component of $\text{FS}_m(\text{Cycle}_n, X)$ containing $[\sigma_{[\alpha]_{\simeq}}]_{\equiv}$. Then*

$$H_{[\alpha]_{\simeq}}, \dots, (\varphi^*)^{\delta_\alpha-1}(H_{[\alpha]_{\simeq}}),$$

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

$$\text{FS}_m(\text{Cycle}_n, X) = \bigoplus_{[\alpha]_{\simeq} \in \text{Acyc}(\overline{X'})/\simeq} \bigoplus_{k=0}^{\delta_\alpha-1} (\varphi^*)^k(H_{[\alpha]_{\simeq}}).$$

Before setting up the proof of the Theorem 5.10, we first present Proposition 5.11. The proof of Proposition 5.11 introduces the new element that distinguishes the proof of Theorem 5.10 from the proofs of Theorems 5.2 and 5.3, and the proofs from [6].

Proposition 5.11. *Let Z_t be one of the connected components of $\overline{X'}$ and suppose Z_t has q vertices. Let $\sigma : [q] \rightarrow V(Z_t)$ be an arrangement and suppose there exists a sequence Σ of $(\text{Cycle}_q, \overline{Z_t})$ -friendly swaps transforming σ into $\rho \circ \sigma \circ \varphi_q^{a_t}$, for some $\rho \in \mathfrak{S}_X$ and $a_t \in \mathbb{Z}_{>0}$. If $\alpha = \alpha_{Z_t}(\sigma)$ is the orientation induced by σ , we have $\pi_\alpha | a_t$.²*

Proof. Since Z_t are connected components of $\overline{X'}$, they form a coarser partition of $V(X')$ than the cliques S_v corresponding to vertices $v \in V(X)$. Hence, we may say that $V(Z_t)$ is a union of several cliques S_v , for v in some set $X_t \subseteq V(X)$. Hence, we will think of Z_t as the complement of the lift of X_t .

The main idea of the proof is to find an orientation $\tilde{\sigma}$, which differs from σ by a sequence of $(\text{Cycle}_q, \overline{Z_t})$ -friendly swaps, and whose period $\pi_{\tilde{\sigma}}$ divides a_t . Once this is established, a simple argument will complete the proof. Before explaining how exactly we find such $\tilde{\sigma}$, let us say a few words about the combinatorial interpretation of the sequence of swaps transforming σ into $\rho \circ \sigma \circ \varphi_q^{a_t}$.

Recall that the arrangement σ is a bijection $\sigma : [q] \rightarrow V(Z_t)$. Hence, we can think of σ as labeling the graph Cycle_q , such that every position $p \in [q]$ on the cycle receives a label $u = \sigma(p) \in Z_t$. Since the condition imposed on Σ is invariant under permuting the labels in the same clique S_v , we may assume that none of the swaps includes two elements of the same clique. In other words, we may assume that the cyclic order of elements of the same clique S_v is unchanged.

For an arbitrary position p , we define the weight of p , denoted by $w(p)$, as the number of times the label $\sigma(p)$ moves counterclockwise in a swap minus the number of times it moves clockwise in Σ . This notion was originally

²Proposition 5.11 implies that $\pi_\alpha | \pi_\sigma$ for all $\sigma \in \mathcal{L}(\alpha)$, since σ is permutation equivalent to $\sigma \circ \varphi_q^{\pi_\sigma}$ by definition and Σ can be taken to be empty. Although this fact is not directly related to the notion of friends-and-strangers graphs, we do not see an easy way to prove it without referencing Proposition 5.11.

introduced by Defant and Kravitz in the proof of Proposition 4.4 of [6]. Since every swap involves one label moving clockwise and one label moving counterclockwise, we must have $\sum_{p \in [q]} w(p) = 0$. The notion of weight turns out to be very useful, because of its relation to adjacent pairs of vertices in $E(Z_t)$.

More precisely, suppose that $p_1, p_2 \in [q]$ satisfy $w(p_1) \geq w(p_2) + d(p_1, p_2)$, where $d(p_1, p_2)$ denotes the counterclockwise distance from p_1 to p_2 on Cycle_q . It is not hard to see that the labels $\sigma(p_1), \sigma(p_2)$ must have participated in the same $(\text{Cycle}_q, \overline{Z_t})$ -friendly swap, implying that $\sigma(p_1)\sigma(p_2) \notin E(Z_t)$. It is natural to say that the positions p_1 and p_2 have *intersecting trajectories* if either $w(p_1) \geq w(p_2) + d(p_1, p_2)$ or $w(p_2) \geq w(p_1) + d(p_2, p_1)$ holds.

As stated earlier, the strategy of the proof is to find $\tilde{\sigma}$ by performing a sequence of $(\text{Cycle}_q, \overline{Z_t})$ -friendly swaps on σ . First of all, let us note that $\tilde{\sigma}$ can still be transformed into $\tilde{\rho} \circ \tilde{\sigma} \circ \varphi_q^{a_t}$, for some $\tilde{\rho} \in \mathfrak{S}_X$, using a sequence $\tilde{\Sigma}$ of $(\text{Cycle}_q, \overline{Z_t})$ -friendly swaps. If $\tilde{\sigma} = \sigma \circ (p \ p + 1)$, then it is not hard to see that the sequence $\tilde{\Sigma} = (p \ p + 1)\Sigma(p + a_t \ p + a_t + 1)$ satisfies the requirements. Hence, the notion of intersecting trajectories carries over to any $\tilde{\sigma}$ obtained from σ through a sequence of $(\text{Cycle}_q, \overline{Z_t})$ -friendly swaps.

Using the notion of the intersecting trajectories, we can finally state precisely and formally our strategy for finding $\tilde{\sigma}$. Namely, the first step of the proof will be showing that one can transform σ into $\tilde{\sigma}$ using a sequence of $(\text{Cycle}_q, \overline{Z_t})$ -friendly swaps such that $\tilde{\sigma}$ satisfies the same conditions as σ and such that no two positions have intersecting trajectories in $\tilde{\sigma}$. Then, we show that for every $\tilde{\sigma}$ is chosen in this way we have $\pi_{\tilde{\sigma}}|_{a_t}$.

Before we present the details of how these two steps are accomplished, let us introduce the permutation $f_\sigma : [q] \rightarrow [q]$ given by $f_\sigma = \sigma^{-1} \circ \rho^{-1} \circ \sigma$. Intuitively, $f_\sigma(p) + a_t$ indicates the position at which the label $\sigma(p)$ ends up at after the sequence of swaps Σ . When σ is clear from the context, or when we are working with only one arrangement, we will write f instead of f_σ . The permutation f splits $[q]$ into orbits, or cycles, and all elements of an orbit are in the same clique of $\overline{Z_t}$. Let us denote the orbit of p by $O_p = \{p, f(p), f^2(p), \dots\}$. Next, we show that all orbits have the same size.

Lemma 5.12. *For any two positions $p_1, p_2 \in [q]$, the orbits O_{p_1} and O_{p_2} have the same size.*

Proof. Since Z_t is connected, it suffices to show the statement is true when $\sigma(p_1)$ and $\sigma(p_2)$ are adjacent in Z_t . Suppose this is indeed the case and let $O_{p_1} = \{p_1, f(p_1), \dots\}$ and $O_{p_2} = \{p_2, f(p_2), \dots\}$ be their orbits. Without loss of generality, we may choose p_1 and p_2 from their orbits such that no elements of O_{p_1} or O_{p_2} lie in (p_1, p_2) , where (p_1, p_2) denotes the open counterclockwise from p_1 to p_2 on Cycle_q .

This choice also implies $(O_{p_1} \cup O_{p_2}) \cap (f(p_1), f(p_2)) = \emptyset$. To see why, suppose there is another element of O_{p_1} , say $f^i(p_1)$, in $(f(p_1), f(p_2))$. Since $\sigma(p_1)$ and $\sigma(f^{i-1}(p_1))$ are in the same clique of $\overline{Z_t}$, $\sigma(f^{i-1}(p_1))$ is adjacent to $\sigma(p_2)$ in Z_t , meaning that the trajectories of p_2 and $f^{i-1}(p_1)$ do not intersect. Similarly, by assumption that the labels of the same clique do not change the cyclic order, the trajectories of p_1 and $f^{i-1}(p_1)$ do not intersect. Since $f^i(p_1)$ is in the interval $(f(p_1), f(p_2))$, $f^{i-1}(p_1)$ must also be in the interval (p_1, p_2) . As this is assumed not the be the case, we conclude $(O_{p_1} \cup O_{p_2}) \cap (f(p_1), f(p_2)) = \emptyset$.

This argument inductively extends to show that $(O_{p_1} \cup O_{p_2}) \cap (f^i(p_1), f^i(p_2)) = \emptyset$, for all $i \geq 0$. In particular, plugging in $i = |O_{p_1}|$ lets us conclude that $f^{|O_{p_1}|}(p_2) = p_2$, meaning $|O_{p_2}|$ divides $|O_{p_1}|$. By symmetry, we also have that $|O_{p_1}|$ divides $|O_{p_2}|$, implying $|O_{p_1}| = |O_{p_2}|$. \square




FIGURE 6. Graphical representation of the orbits of f , along with the trajectories of $p_1, p_2, f(p_1), f(p_2)$. The positions in the top row represent the arrangement σ before performing the sequence of $(\text{Cycle}_q, \overline{Z_t})$ -friendly swaps, while the bottom row represents the positions in the permutation after this sequence is performed. The arrows in the diagram represent schematically how labels move throughout this process. Note that in this picture, the trajectories of p_1 and p_2 do not intersect, while the trajectories of $f(p_1)$ and $f(p_2)$ intersect.

Lemma 5.13. *For any $p \in [q]$ we have $\sum_{k=0}^{|O_p|-1} w(f^k(p)) = 0$.*

Proof. In fact, we will show that for arbitrary $p_1, p_2 \in [q]$ with $\sigma(p_1)\sigma(p_2) \in E(Z_t)$ we have

$$(10) \quad \sum_{k=0}^{|O_{p_1}|-1} w(f^k(p)) = \sum_{k=0}^{|O_{p_1}|-1} w(f^k(q)).$$

Since Z_t is connected, relation (10) suffices to show that the sums of weights in all orbits are the same. Because the sum of all weights is 0, the sum of weights in each orbit must also be 0.

To show (10), we note that since the trajectories of $f^k(p_1)$ and $f^k(p_2)$ do not intersect, we can write $w(f^k(p_1)) = w(f^k(p_2)) + d(f^k(p_1), f^k(p_2)) - d(f^{k+1}(p_1), f^{k+1}(p_2))$. Summing over all $k \in \{0, \dots, |O_{p_1}| - 1\}$, the summands $d(f^k(p_1), f^k(p_2))$ telescope and we obtain (10). \square

Lemma 5.14. *Suppose that the trajectories of p_1 and p_2 intersect. Then, there exists some $k \in \{1, \dots, |O_{p_1}| - 1\}$ for which the trajectories of $f^k(p_1)$ and $f^k(p_2)$ also intersect.*

Proof. Since the trajectories of p_1 and p_2 intersect, without loss of generality we may suppose that $w(p_1) \geq w(p_2) + d(p_1, p_2)$. Hence, equation (10) gives

$$w(f(p_1)) + \dots + w(f^{|O_{p_1}|-1}(p_1)) + d(p_1, p_2) \leq w(f(p_2)) + \dots + w(f^{|O_{p_1}|-1}(p_2)).$$

Assuming the trajectories of $f^k(p_1)$ and $f^k(p_2)$ do not intersect for $k \in \{1, \dots, |O_{p_1}| - 1\}$, we deduce that $w(f^k(p_1)) = w(f^k(p_2)) + d(f^k(p_1), f^k(p_2)) - d(f^{k+1}(p_1), f^{k+1}(p_2))$. Summing over $k \in \{1, \dots, |O_{p_1}| - 1\}$ gives

$$w(f(p_1)) + \dots + w(f^{|O_{p_1}|-1}(p_1)) = w(f(p_2)) + \dots + w(f^{|O_{p_1}|-1}(p_2)) + d(f(p_1), f(p_2)) - d(p_1, p_2).$$

Combined with the previous inequality, this implies $d(f(p), f(q)) \leq 0$, which is absurd. The conclusion is that there exists another index $k \in \{1, \dots, |O_p| - 1\}$ for which the trajectories of $f^k(p)$ and $f^k(q)$ intersect. \square

Lemma 5.15. *There exists an arrangement $\tilde{\sigma}$, differing from σ by a sequence of $(\text{Cycle}_q, \overline{Z_t})$ -friendly swaps, which has no intersecting trajectories.*

Proof. Suppose that $\tilde{\sigma}$ is the arrangement with the minimal number of intersecting trajectories among all arrangements obtainable from σ through a sequence of $(\text{Cycle}_q, \overline{Z_t})$ -friendly swaps. Under the assumption that $\tilde{\sigma}$ has a pair of intersecting trajectories, we will show how to reduce the number of pairs having intersecting trajectories.

Before presenting the algorithm which accomplishes this, let us briefly discuss the effects of a $(\text{Cycle}_q, \overline{Z_t})$ -friendly swap on the orbit structure of f_σ and intersecting trajectories. Suppose that $\sigma_1 = \sigma \circ (p \ p+1)$ is a $(\text{Cycle}_q, \overline{Z_t})$ -friendly swap, and that the trajectories of $p, p+1$ were intersecting for σ . Then, these trajectories are not intersecting in σ_1 . Moreover, if the trajectories of $f_\sigma^{-1}(p), f_\sigma^{-1}(p+1)$ are intersecting for σ , they are not intersecting for σ_1 and vice versa. This shows that the number of intersecting pairs does not increase after performing a $(\text{Cycle}_q, \overline{Z_t})$ -friendly swap on a pair with intersecting trajectories. Moreover, the number of intersecting pairs decreases if the pair $f_\sigma^{-1}(p), f_\sigma^{-1}(p+1)$ is also intersecting in σ . These changes are visually presented in Figure 7.




FIGURE 7. For a given arrangement σ , the left diagram represents the trajectories of σ , while the right one represents the trajectories of $\sigma_1 = \sigma \circ (p \ p+1)$. The letters u, v, u', v' represent labels assigned to relevant positions before and after the sequence of swaps is performed, under the convention that $u = \sigma(p), v = \sigma(p+1), u' = \sigma(f_\sigma^{-1}(p)), v' = \sigma(f_\sigma^{-1}(p+1))$.

Suppose now that $\tilde{\sigma}$ contain an intersecting pair of positions, say p_1, p_2 . By Lemma 5.14 the trajectories of $f^k(p)$ and $f^k(q)$ intersect for some $k \in \{1, \dots, |O_p| - 1\}$. Let us consider the smallest such k .

The main goal of the sequence of swaps we will construct is to obtain an arrangement σ' in which $f^k(p_1)$ and $f^k(p_2)$ are adjacent, and then perform a $(\text{Cycle}_q, \overline{Z_t})$ -friendly swap on these two positions. In case $k = 1$, by the

above discussion, this will reduce the intersection number, which is exactly what we want. For $k > 1$, this will reduce to the case when the trajectories of $f^{k-1}(p_1)$ and $f^{k-1}(p_2)$ intersect, which can be dealt with by repeating the above procedure.

Let us now explain how we make $f^k(p_1), f^k(p_2)$ adjacent. It is simple to see that if the trajectories of a and b intersect, then for every $c \in (a, b)$ the trajectory of c intersects that of either a or b . Hence, for every $x \in (f^k(p_2), f^k(p_1))$, we either have that the trajectory of w intersects the trajectory of $f^k(p_1)$ or the trajectory of $f^k(p_2)$. For $i \in \{1, 2\}$, let S_i be the set of those x whose trajectories intersect the trajectory of $f^k(p_i)$. By applying the above argument again, we see that the trajectory of $s_1 \in S_1$ intersects the trajectory of every $s_2 \in S_2 \cap (s_1, f^k(p_1))$. Hence, every such pair (s_1, s_2) has labels which are adjacent in $\overline{\mathbb{Z}}_t$. Therefore one may perform a sequence of $(\text{Cycle}_q, \overline{\mathbb{Z}}_t)$ -friendly swaps to obtain an arrangement in which all positions which have intersecting trajectories with $f^k(p_2)$ come before those positions which have intersecting trajectories with $f^k(p_1)$, in the interval $(f^k(p_2), f^k(p_1))$. Then, we may perform a sequence of $(\text{Cycle}_q, \overline{\mathbb{Z}}_t)$ -friendly swaps which makes $f^k(p_2), f^k(p_1)$ adjacent. As explained above, this suffices to complete the proof. \square




FIGURE 8. Schematic representation of the proof of Lemma 5.15. The left diagram illustrates the fact that if the trajectories of a and b intersect and $c \in (a, b)$, then the trajectory of c intersects the trajectory of either a or b . The middle and the right diagram represent steps in transforming $\tilde{\sigma}$ into an arrangement in which $f^k(p_1)$ and $f^k(p_2)$ are adjacent.

Lemma 5.16. *If $\tilde{\sigma}$ is an arrangement in which no pair of positions has intersecting trajectories, then $\pi_{\tilde{\sigma}}|_{a_t}$.*

Proof. We begin by showing that $w(p)$ is the same for all $p \in [q]$. Let p and $p+1$ be two adjacent vertices of Cycle_q . Since their trajectories do not intersect, we have $w(p) \leq w(p+1)$ and $w(p+1) \leq w(p) + (q-1)$. In case $w(p) \neq w(p+1)$, we must have $w(p+1) \geq w(p) + 1$. It is not hard to see that, in this case, the trajectory of $f^{-1}(f(p)+1)$ must intersect either the trajectory of p or the trajectory of $p+1$. As $\tilde{\sigma}$ has no pairs with intersecting trajectories, this is impossible and therefore $w(p) = w(p+1)$. As this holds for all positions $p \in [q]$, we conclude that $w(p)$ is the same for all $p \in [q]$. Since $\sum_{p \in [q]} w(p) = 0$, we have $w(p) = 0$ for all $p \in [q]$.




FIGURE 9. Graphical representation of the fact that the trajectory of $f^{-1}(f(p)+1)$ intersects either the trajectory of p or the trajectory of $p+1$ in the proof of Lemma 5.16.

This means that $f(p) = p$ for all $p \in [q]$, and hence the empty sequence of $(\text{Cycle}_q, \overline{\mathbb{Z}}_t)$ -friendly swap transforms $\tilde{\sigma}$ into $\rho \circ \tilde{\sigma} \circ \varphi_q^{a_t}$. In other words, $\tilde{\sigma} \equiv \tilde{\sigma} \circ \varphi_q^{a_t}$, implying $\pi_{\tilde{\sigma}}|_{a_t}$, which completes the proof of Lemma 5.16. \square

Combining Lemmas 5.15 and 5.16 shows that there exists an arrangement $\tilde{\sigma}$, differing from σ by a sequence of $(C_q, \overline{\mathbb{Z}}_t)$ -friendly swaps, and having $\pi_{\tilde{\sigma}}|_{a_t}$. Let us now argue why this suffices to conclude $\pi_{\alpha}|_{a_t}$, and assume, towards contradiction, that this is not the case.

We have argued previously that replacing σ with any arrangement τ differing from it by a sequence of $(C_q, \overline{\mathbb{Z}}_t)$ -friendly swaps does not change the property that τ can be transformed into $\rho' \circ \tau \circ \varphi_q^{a_t}$ through a sequence of

(C_q, \overline{Z}_t) -friendly swaps, for some $\rho' \in \mathfrak{S}_X$. Let $\tau \in \mathcal{L}(\alpha)$ be the arrangement with the minimal period $\pi_\tau = \pi_\alpha$. Since $\tau \circ \varphi_q^{a_t} = \tau \circ \varphi_q^{a_t - m\pi_\alpha}$ for any $m \geq 0$, we may assume that $1 \leq a_t < \pi_\alpha$ and that τ can be transformed into $\rho' \circ \tau \circ \varphi_q^{a_t}$ through a sequence of (C_q, \overline{Z}_t) -friendly swaps. However, since $\tilde{\sigma}$ differs from τ by a sequence of (C_q, \overline{Z}_t) -friendly swaps, we conclude that $\tilde{\sigma}$ also satisfies the same property as τ with $a_t < \pi_\alpha$. This presents a contradiction, since $\pi_{\tilde{\sigma}}|_{a_t}$ and $\pi_{\tilde{\sigma}} \geq \pi_\alpha$ by definition of π_α . This contradiction completes the proof of Proposition 5.11. \square

Equipped with Proposition 5.11, we are ready to prove Theorem 5.10.

Proof of Theorem 5.10. We structure the proof in a similar way to the proofs of Theorems 5.2 and 5.3. However, the analysis of the arising graph G is more involved than earlier and uses Proposition 5.11 in an important way.

Let G be the graph arising from $\text{FS}(\text{Cycle}_n, X')$ after adding the set of edges $E_0 = \{\sigma\tau : \sigma \equiv \tau\}$. Corollary 2.10 states that the connected components of the graph $\text{FS}_m(\text{Cycle}_n, X)$ are the images of the connected components of G under the projection $\sigma \mapsto [\sigma]_{\equiv}$. Hence, the main focus of the proof is to describe the connected components of G . More precisely, if $L_{[\alpha]_{\simeq}} = G|_{\{\sigma : [\sigma]_{\equiv} \in V(H_{[\alpha]_{\simeq}})\}}$ is the preimage of $H_{[\alpha]_{\simeq}}$ under the discussed projection, we aim to show that

$$(11) \quad G = \bigoplus_{\alpha \in \text{Acyc}(\overline{X'})/\simeq} \bigoplus_{i=0}^{\delta_\alpha-1} (\varphi^*)^i L_{[\alpha]_{\simeq}}.$$

Since the projection operation commutes with the automorphism φ^* , which is to say that the projection of $(\varphi^*)^i L_{[\alpha]_{\simeq}}$ is $(\varphi^*)^i H_{[\alpha]_{\simeq}}$, showing (11) would suffice to prove the second part of the theorem. The fact that $H_{[\alpha]_{\simeq}}, \dots, (\varphi^*)^{\delta_\alpha-1} H_{[\alpha]_{\simeq}}$ are distinct isomorphic connected components follows directly, since φ^* is an automorphism of the graph $\text{FS}_m(\text{Cycle}_n, X)$.

Let us begin the proof of (11) by noting that there are no edges of G out of the set $\mathcal{L}([\alpha]_{\simeq})$. To see why, we pick an arbitrary edge $\sigma\tau \in E(G)$ and show $\alpha_{\overline{X'}}(\sigma) \simeq \alpha_{\overline{X'}}(\tau)$. If $\sigma\tau \in E_0$, we must have $\sigma \equiv \tau$, and hence $\alpha_{\overline{X'}}(\sigma) \equiv \alpha_{\overline{X'}}(\tau)$. On the other hand, if $\sigma\tau \in E(\text{FS}(\text{Cycle}_n, X'))$, we have $\alpha_{\overline{X'}}(\sigma) \approx \alpha_{\overline{X'}}(\tau)$, and therefore also $\alpha_{\overline{X'}}(\sigma) \sim \alpha_{\overline{X'}}(\tau)$. Hence, all edges of G must be fully contained within one of the sets $\mathcal{L}([\alpha]_{\simeq})$, which gives the following decomposition of G :

$$(12) \quad G = \bigoplus_{\alpha \in \text{Acyc}(\overline{X'})/\simeq} G|_{\mathcal{L}([\alpha]_{\simeq})}.$$

In order to infer (11) from (12), we need to show that

$$(13) \quad G|_{\mathcal{L}([\alpha]_{\simeq})} = \bigoplus_{i=0}^{\delta_\alpha-1} (\varphi^*)^i L_{[\alpha]_{\simeq}}.$$

Since the equivalence relation \simeq is coarser than \sim , the equivalence class $[\alpha]_{\simeq}$ can be written as a disjoint union of equivalence classes $[\alpha_1]_{\sim}, \dots, [\alpha_k]_{\sim}$. Proposition 5.9 implies that

$$(14) \quad \text{FS}(\text{Cycle}_n, X')|_{\mathcal{L}([\alpha]_{\simeq})} = \bigoplus_{j=1}^k \bigoplus_{i=0}^{\nu-1} (\varphi^*)^i J_{[\alpha_k]_{\sim}},$$

where $J_{[\alpha_k]_{\sim}}$ is the connected component of $\text{FS}(\text{Cycle}_n, X')$ containing an arbitrarily chosen vertex $\sigma_{[\alpha_k]_{\sim}} \in \mathcal{L}([\alpha_k]_{\sim})$. Since $\alpha_1, \dots, \alpha_k$ are all equivalent under \simeq , we may choose $\sigma_{[\alpha_1]_{\sim}}, \dots, \sigma_{[\alpha_k]_{\sim}}$ to be equivalent under \equiv . We now present two lemmas, which show that there exists an edge of E_0 between components $(\varphi^*)^i J_{[\alpha_l]_{\sim}}$ and $(\varphi^*)^j J_{[\alpha_m]_{\sim}}$ if and only if $\delta_\alpha | i - j$.

Lemma 5.17. *If δ_α divides $i - j$, then there exists an edge of E_0 between components $(\varphi^*)^i J_{[\alpha_l]_{\sim}}$ and $(\varphi^*)^j J_{[\alpha_m]_{\sim}}$.*

Proof. Let us first consider the case $l = m$, as the case $l \neq m$ can be easily deduced from it. The main idea is to show that, for every arrangement $\sigma \in (\varphi^*)^i J_{[\alpha_l]_{\sim}}$, there exists an edge of E_0 between σ and a vertex of $(\varphi^*)^{i+\pi_t} J_{[\alpha_l]_{\sim}}$, where π_t is the period of an induced orientation $\alpha^{(t)}$ on Z_t .

Let β be the orientation induced by σ on $\overline{X'}$, and let $\beta^{(t)}$ be its restriction to Z_t . Because $\sigma \in \mathcal{L}([\alpha]_{\simeq})$, and consequently $\beta^{(t)} \simeq \alpha^{(t)}$, the period of $\beta^{(t)}$ must be π_t .

Since the period of $\beta^{(t)}$ is π_t , there exists a sequence of π_t positive single flips that transforms $\beta^{(t)}$ into $\beta_1^{(t)}$ which is permutation equivalent to $\beta^{(t)}$. Since there are no edges of $\overline{X'}$ leaving Z_t , the same sequence of π_t flips can be performed on β to get a orientation β_1 with $\beta_1 \equiv \beta$. In particular, this means that there exists a linear extension $\sigma_1 \in \mathcal{L}(\beta_1)$ for which $\sigma \equiv \sigma_1$. The goal is now to show $\sigma_1 \in (\varphi^*)^{i+\pi_t} J_{[\alpha_l]_{\sim}}$.

Let β_2 be the orientation induced on \overline{X} by $\sigma \circ \varphi^{\pi_t}$. Note that β_2 can also be obtained from β through a sequence of π_t positive single flips. Using Corollary 5.1, this implies $\beta_1 \approx \beta_2$.

Since $\sigma \circ \varphi^{\pi_t} \in (\varphi^*)^{\pi_t+i} J_{[\alpha_l]_{\sim}}$, Theorem 4.1 of [6] shows that the vertex σ_1 also lies in $(\varphi^*)^{\pi_t+i} J_{[\alpha_l]_{\sim}}$. Since $\sigma_1 \equiv \sigma$, we see there exists an edge of E_0 between σ and a vertex of $(\varphi^*)^{\pi_t+i} J_{[\alpha_l]_{\sim}}$.

By repeating the above argument and using the transitivity of the relation \equiv , we conclude that there exists an edge of E_0 between σ and a vertex of $(\varphi^*)^{i+\sum_{t=1}^r \lambda_t \pi_t} J_{[\alpha_l]_{\sim}}$, where $\lambda_1, \dots, \lambda_r$ are arbitrary integer coefficients. Using Bezout's theorem, we conclude there exist coefficients $\lambda_1, \dots, \lambda_t$ for which $\sum_{t=1}^r \lambda_t \pi_t = u \gcd(\pi_1, \dots, \pi_r) = u \delta_\alpha$, where u is an arbitrary integer. Hence, if δ_α divides $i - j$, we have an edge between components every vertex of $(\varphi^*)^i J_{[\alpha_l]_{\sim}}$ and $(\varphi^*)^j J_{[\alpha_l]_{\sim}}$. Since $\sigma_{[\alpha_1]_{\sim}}, \dots, \sigma_{[\alpha_k]_{\sim}}$ are all permutation equivalent, we have an edge of E_0 between $(\varphi^*)^i J_{[\alpha_l]_{\sim}}$ and $(\varphi^*)^j J_{[\alpha_m]_{\sim}}$, for all $l, m \in [k]$. As before, since every vertex of $(\varphi^*)^i J_{[\alpha_m]_{\sim}}$ has an edge of E_0 to $(\varphi^*)^j J_{[\alpha_m]_{\sim}}$, when $\delta_\alpha | i - j$, by transitivity of \equiv we conclude that there exists an edge of E_0 between $(\varphi^*)^i J_{[\alpha_l]_{\sim}}$ and $(\varphi^*)^j J_{[\alpha_m]_{\sim}}$ whenever δ_α divides $i - j$. This completes the proof of Lemma 5.17. \square

Lemma 5.18. *If there exists an edge of E_0 between components $(\varphi^*)^i J_{[\alpha_l]_{\sim}}$ and $(\varphi^*)^j J_{[\alpha_m]_{\sim}}$, then δ_α divides $i - j$.*

Proof of Lemma 5.18. As in the proof of Lemma 5.17, we begin by focusing on the case $l = m$, and then reduce the general case to this one. Our argument has two main steps – in the first step, we separate the problem into smaller subproblems, by restricting it to one of the connected components Z_t . Then, we apply Proposition 5.11 to solve each of the subproblems.

Suppose there exists an edge of E_0 between $\sigma \in \varphi^i J_{[\alpha_l]_{\sim}}$ and $\rho \circ \sigma \in \varphi^j J_{[\alpha_l]_{\sim}}$, for some $\rho \in \mathfrak{S}_X$. In other words, there exists a sequence of $a = j - i$ positive single flips and b double flips that transforms $\beta = \alpha_{\overline{X}}(\sigma)$ into a permutation equivalent orientation $\rho(\beta) = \alpha_{\overline{X}}(\rho \circ \sigma)$.

The first step is to associate every flip to the component Z_t in which it is performed. More precisely, if the sequence of flips Σ transforming σ into $\rho \circ \sigma$ consists of single flips at vertices u_1, \dots, u_a , and double flips at pairs $(v_1, w_1), \dots, (v_b, w_b)$, we consider three different types of flips: single flips, double flips with both v_i and w_i in the same connected component Z_t , and double flips with v_i and w_i in different connected components. We split the flips of Σ into r subsequences of flips, denoted $\Sigma_1, \dots, \Sigma_r$, so that Σ_t contains only flips at vertices of Z_t . Hence, the first two types of flips are simply associated to the component they are performed in. For each flip of the third type (v_i, w_i) , we split it into a positive single flip at v_i and a negative single flip at w_i , which are then placed into subsequences corresponding to their respective components. The result of this process is splitting Σ into $\Sigma_1, \dots, \Sigma_r$, where Σ_t is a legal sequence of single and double flips performed on the restriction $\beta|_{Z_t}$. Moreover, the result of applying Σ_t to $\beta^{(t)} = \beta|_{Z_t}$ is precisely the orientation $\rho(\beta)|_{Z_t}$.

Let $a_{t,+}$ ($a_{t,-}$, resp.) denote the number of positive (negative, resp.) single flips in Σ_t , and let $a_t = a_{t,+} - a_{t,-}$ denote their difference. Because the number of positive flips in Σ is a , and since splitting every double flip amounts to adding one positive and one negative single flip, we conclude that $a_1 + \dots + a_r = a$. The main goal now is to show that $\pi_t | a_t$, which will immediately imply $\delta_\alpha | a$.

Let us now rephrase the problem in the setting of a single connected component Z_t . As discussed earlier, by Corollary 4.2 of [6], Σ_t is equivalent to a sequence of a_t positive flips (or $-a_t$ negative flips, if $a_t < 0$). The problem now simplifies to showing that, if $\beta^{(t)}, \rho(\beta^{(t)}) \in \text{Acyc}(Z_t)$ are permutation equivalent orientations and, and $\rho(\beta^{(t)})$ is obtained from $\beta^{(t)}$ after a sequence of a_t positive single flips and some double flips, we must have $\pi_t | a_t$. This is precisely what Proposition 5.11 states, albeit in slightly different terminology.

If we denote $|V(Z_t)|$ by q , Corollary 4.2 of [6] shows that $\rho(\beta^{(t)})$ must be double flip equivalent to the acyclic orientation induced by $\sigma^{(t)} \circ \varphi_q^{a_t}$, where we think of the arrangement $\sigma^{(t)}$ as mapping Cycle_q to Z_t . In other words, the orientations induced by $\sigma^{(t)}$ and $\rho^{-1} \circ \sigma^{(t)} \circ \varphi_q^{a_t}$ are double flip equivalent. Proposition 5.9 implies that these two arrangement differ by a sequence of $(\text{Cycle}_q, \overline{Z_t})$ -friendly swaps. In this context, Proposition 5.11 shows that $\pi_t | a_t$, which is exactly what we need to complete the proof in the case $l = m$.

Suppose now that there exists an edge of E_0 between vertices $\sigma \in (\varphi^*)^i J_{[\alpha_l]_{\sim}}$ and $\tau \in (\varphi^*)^j J_{[\alpha_m]_{\sim}}$. Since $\sigma_{[\alpha_l]_{\sim}} \equiv \sigma_{[\alpha_m]_{\sim}}$, we know that there exists an edge from every vertex of $\rho \circ \tau \in (\varphi^*)^j J_{[\alpha_l]_{\sim}}$ to a vertex of $\tau \in (\varphi^*)^j J_{[\alpha_m]_{\sim}}$. Hence, by transitivity of \equiv , we conclude there exists an edge of E_0 between $\sigma \in (\varphi^*)^i J_{[\alpha_l]_{\sim}}$ and $\rho \circ \tau \in (\varphi^*)^j J_{[\alpha_l]_{\sim}}$, which is possible only if $\delta_\alpha | j - i$, as desired. \square

Combining Lemmas 5.17 and 5.18 with (14), we conclude that $G|_{\mathcal{L}([\alpha]_{\sim})}$ contains exactly δ_α connected components, with the vertex sets corresponding to $\bigcup_{j=1}^k \bigcup_{i \equiv i_0 \pmod{\delta_\alpha}} V((\varphi^*)^i J_{[\alpha_k]_{\sim}})$. Moreover, these δ_α connected components differ only by an application of φ^* .

Since $L_{[\alpha]_{\sim}}$ is one of these δ_α connected components, decomposition (13) now becomes obvious. As discussed above, this suffices to infer (11), and to complete the proof of Theorem 5.10. \square

Corollary 5.19. *The graph $\text{FS}_m(\text{Cycle}_n, X)$ is connected if and only if the complement of the lift of X is a forest of trees of coprime sizes.*

Proof. Note that if X' is the complement of a forest of trees of coprime sizes, by the Corollary 4.14 of [6] we have that $\text{FS}(\text{Cycle}_n, X')$ is connected. Hence, Corollary 2.9 implies that $\text{FS}_m(\text{Cycle}_n, X)$ must be connected. Establishing the other direction is somewhat trickier, and requires the machinery developed in this section.

By Theorem 5.10, if the graph $\text{FS}_m(\text{Cycle}_n, X)$ has only one connected component, we must have $|\text{Acyc}(\overline{X'})| = 1$ and $\delta_\alpha = 1$ for the unique element $[\alpha]_{\simeq} \in \text{Acyc}(\overline{X'})/\simeq$. Unlike the proof of Corollary 4.14 in [6], which uses a systematic way to count the number of toric acyclic orientations of a graph using its Tutte polynomial, we present a direct and straightforward argument characterizing all graphs with a single orientation up to \simeq equivalence.

Let us begin by showing the complement of X has no cycles. If there was a cycle in \overline{X} , consisting of vertices v_1, \dots, v_k in this order, consider the orientation $\alpha \in \text{Acyc}(\overline{X'})$ which directs all edges from $u \in S_{v_i}$ to $w \in S_{v_j}$ whenever $i < j$. If we imagine S_{v_1}, \dots, S_{v_k} are placed in the plane in the counterclockwise order, it is not hard to see that every cycle of $\overline{X'}$, having one vertex in each of the cliques S_{v_i} , has exactly one edge oriented clockwise. For each fixed cycle of this type, it is not hard to see that even after several flips are performed, there is still exactly one edge oriented clockwise. We have thus shown that this property is maintained for every orientation α' obtained from α through a sequence of flips. However, there exist orientations of $\overline{X'}$ which do not have this property, meaning that $\text{Acyc}(\overline{X'})/\simeq$ has more than one element, which is impossible. An example of such orientations is an orientation in which all edges of α are reversed.

Since \overline{X} has no cycles, it is a forest. Suppose $v \in X$ is a vertex of capacity $c \geq 2$. We will show that it can be adjacent to a single vertex of capacity 1. For every vertex u adjacent to S_v , we let $p(u, \alpha)$ be the number of edges directed from u to S_v under α . We note that throughout a sequence of flips the quantity $p(u_1, \alpha) - p(u_2, \alpha)$ remains constant modulo c , for any vertices u_1, u_2 in the neighborhood of S_v . If S_v is adjacent to at least two vertices, we consider the orientation α in which all edges incident to S_v are oriented away from S_v and an orientation α' in which exactly one edge incident to S_v is reversed. By the above remark, α and α' do not differ by a sequence of flips as long as there are at least 2 neighbors of S_v in $\overline{X'}$. Again, this is impossible since $\text{Acyc}(\overline{X'})/\simeq$ is assumed to have only one equivalence class.

We conclude $\overline{X'}$ is a forest, consisting of trees T_1, \dots, T_m . Since $\text{FS}_m(\text{Cycle}_n, X)$ is connected, every $\alpha \in \text{Acyc}(\overline{X'})/\simeq$ has $\delta_\alpha = 1$. Recall that δ_α is defined as the greatest common divisor of periods of orientations $\alpha^{(i)}$, where $\alpha^{(i)}$ is the restriction of α to T_i . We will show now that the period of $\alpha^{(i)}$ is $|T_i|$. Since, up to toric equivalence, there is only one acyclic orientation of $\overline{X'}$, we may choose a root u_i in every tree T_i and assume that $\alpha^{(i)}$ directs the edges of T_i away from the root. For this specific orientation, it is clear that at least $|T_i|$ positive flips are needed in order to transform it back into α . Hence, the period of $\alpha^{(i)}$ is $|T_i|$. Since $\delta_\alpha = 1$, this implies $\gcd(|T_1|, \dots, |T_m|) = 1$, which completes the proof. \square

6. FURTHER RESEARCH DIRECTIONS

As mentioned in the introduction, Theorem 1.3 determines the correct dependence, up to a factor of $n^{o(1)}$, of the threshold probability $p_{\text{bip}} = p_{\text{bip}}(n)$ for which $\text{FS}(X, Y)$ has exactly two connected components, when $X, Y \sim \mathcal{G}(K_{n,n}, p)$. It would be interesting to close the gap, in both bipartite and non-bipartite settings, and determine the exact threshold.

If one believes that the main obstruction to the connectivity of randomly generated friends-and-strangers graphs are isolated vertices (as in the case of the usual random graphs), then one would conjecture that the threshold probabilities p_{gen} and p_{bip} for the connectivity of friends-and-strangers graphs would be the same as for containing isolated vertices, i.e. $p_{\text{gen}}, p_{\text{bip}} = \left(\frac{\log n}{n}\right)^{1/2}$ (cf. Remark 2.5). Hence, we make the following conjecture, which would determine the threshold probabilities precisely.

Conjecture 6.1. *For $p = \omega\left(\frac{\log^{1/2} n}{n^{1/2}}\right)$ and large n we have the following two statements:*

- *if $X, Y \sim \mathcal{G}(n, p)$, then $\text{FS}(X, Y)$ is connected with high probability.*
- *if $X, Y \sim \mathcal{G}(K_{n,n}, p)$, then $\text{FS}(X, Y)$ has exactly two connected components with high probability.*

In the context of multiplicity graphs, there exists a further generalization of this concept which allows for both graphs X and Y to have multiplicities. More precisely, if X, Y are multiplicity graphs with lists of multiplicities $c^{(X)}, c^{(Y)}$, having the same total capacity, then we can define the double multiplicity friends-and-strangers graph $\text{FS}_{m,m}(X, Y)$ whose vertices are arrangements of labels represented by the vertices of X which are being placed onto the vertices of Y , such that the label $u \in V(X)$ appears exactly $c_u^{(X)}$ times and such that exactly $c_v^{(Y)}$ labels are placed onto every vertex of $v \in V(Y)$. Many of the questions addressed in this paper can be posed for

double-multiplicity friends-and-strangers graphs. For example, since this paper exactly characterizes all multiplicity lists for which $\text{FS}_m(X, \text{Star}_n)$ and $\text{FS}_m(\text{Star}_n, X)$ are connected, it is natural to ask for an characterizations of all multiplicity lists $c^{(X)}, c^{(\text{Star}_n)}$ which make $\text{FS}_{m,m}(X, \text{Star}_n)$ connected. In particular, we believe that the following characterization can be derived.

Conjecture 6.2. *Let X be a connected simple graph, and let $c^{(X)}, c^{(\text{Star}_n)}$ be two multiplicity lists, in which the center of Star_n receives multiplicity $k = c_{\emptyset}^{(\text{Star}_n)}$. Then $\text{FS}_{m,m}(X, \text{Star}_n)$ is connected if and only if X contains no k -bridge in which all vertices have multiplicity 1.*

Finally, this paper mainly investigates connectivity of various models of friends-and-strangers graphs. However, one can still ask questions about higher connectivity of $\text{FS}(X, Y)$ and $\text{FS}_m(X, Y)$. For example, one can pose the following question in case X, Y are random graphs.

Question 6.3. *Let $p(n) = n^{-1/2+o(1)}$ and let X, Y be random graphs from $\mathcal{G}(n, p)$. For which values of k is the graph $\text{FS}(X, Y)$ is k -connected with high probability?*

ACKNOWLEDGEMENTS

This research was conducted at the University of Minnesota Duluth REU, with generous support from Jane Street Capital, the NSA (grant number H98230-22-1-0015), and Princeton University. I would like to thank Joe Gallian for organizing this REU program, Noah Kravitz and Colin Defant for introducing me to this problem, proposing many new research directions, and providing useful feedback on early drafts of this paper. Further, I would like to thank Kiril Bangachev for carefully reading the paper and pointing out several mistakes in its early drafts, as well as Maya Sankar and Mitchell Lee for helpful discussions about the problem, reading drafts of this paper, and generously sharing their ideas with me. Finally, I would like to thank an anonymous reviewer for his helpful comments and pointing out that reference [4] may lead to the proof of Theorem 1.2.

REFERENCES

- [1] N. Alon, C. Defant, and N. Kravitz, Typical and extremal aspects of friends-and-strangers graphs. *Journal of Combinatorial Theory, Series B*, (2022).
- [2] K. Bangachev, On the asymmetric generalizations of two extremal questions on friends-and-strangers graphs. *European Journal of Combinatorics*, **104**, (2022).
- [3] T. Brailovskaya, G. Gowri, S. Yu, and E. Winfree, Reversible computation using swap reactions on a surface. *DNA Computing and Molecular Programming*, (2019), 174–196.
- [4] B. Bollobás, S. Janson and A. Scott, Packing random graphs and hypergraphs, *Random Structures and Algorithms* **51** (2017), 3–13.
- [5] C. Defant, D. Dong, A. Lee, and M. Wei, Connectedness and cycle spaces of friends-and-strangers graphs, *ArXiv preprint, arXiv:2209.01704*, (2022).
- [6] C. Defant and N. Kravitz, Friends and strangers walking on graphs. *Combinatorial Theory*, (2021).
- [7] M. Develin, M. Macauley, and V. Reiner, Toric partial orders. *Trans. Amer. Math. Soc.*, **368** (2016), 2263–2287.
- [8] R. Jeong, On structural aspects of friends-and-strangers graphs, *ArXiv preprint, arXiv:2203.10337*, (2022).
- [9] R. Jeong, Diameters of connected components of friends-and-strangers graphs are not polynomially bounded, *ArXiv preprint, arXiv:2201.00665*, (2022).
- [10] D. Kornhauser, G. Miller, and P. Spirakis, Coordinating pebble motion on graphs, the diameter of permutation groups, and applications. *25th Annual Symposium on Foundations of Computer Science*, (1984).
- [11] C. M. Reidys, Acyclic orientations of random graphs. *Advances in Applied Mathematics*, **21**, (1998), 181–192.
- [12] R. P. Stanley, An equivalence relation on the symmetric group and multiplicity-free flag h -vectors. *Journal of Combinatorics*, **3** (2012), 277–298.
- [13] R. M. Wilson, Graph puzzles, homotopy, and the alternating group. *Journal of Combinatorial Theory, Series B*, **16** (1974), 86–96.
- [14] L. Wang and Y. Chen, Connectivity of friends-and-strangers graphs on random pairs. *ArXiv preprint, arXiv:2208.00801*, (2022).

DEPARTMENT OF MATHEMATICS, PRINCETON UNIVERSITY, PRINCETON, NJ 08540, USA
Email address: `aleksam@princeton.edu`