

ON THE DIAMETERS OF FRIENDS-AND-STRANGERS GRAPHS

Ryan Jeong^{*1}

¹*Department of Statistics, Stanford University, Stanford, California, U.S.A.
rsjeong@stanford.edu*

Submitted: May 27, 2022; Accepted: Apr 3, 2024; Published: Sep 30, 2024

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

Abstract. Given simple graphs X and Y on the same number of vertices, the friends-and-strangers graph $\text{FS}(X, Y)$ has as its vertices all bijections from $V(X)$ to $V(Y)$, where two bijections are adjacent if and only if they differ on two adjacent elements of $V(X)$ with images adjacent in Y . We study the diameters of connected components of friends-and-strangers graphs: the diameter of a component of $\text{FS}(X, Y)$ corresponds to the largest number of swaps necessary to go from one configuration in the component to another. We show that any component of $\text{FS}(\text{Path}_n, Y)$ has $O(n^2)$ diameter and that any component of $\text{FS}(\text{Cycle}_n, Y)$ has $O(n^4)$ diameter, improvable to $O(n^3)$ whenever $\text{FS}(\text{Cycle}_n, Y)$ is connected. Answering a question raised by Alon, Defant, and Kravitz in the negative, we use an explicit construction to show that there exist n -vertex graphs X and Y such that $\text{FS}(X, Y)$ has a component with $e^{\Omega(n)}$ diameter. We conclude with several suggestions for future research.

Keywords. Friends-and-strangers graphs, diameter, extremal combinatorics, lower bounds, paths, cycles, token swapping, interchange process

Mathematics Subject Classifications. 05C12, 05C35, 05C38

1. Introduction

1.1. Background and Motivation

Let X and Y be n -vertex simple graphs. Interpret the vertices of X as positions, and the vertices of Y as people: say two people in the vertex set of Y are friends if they are adjacent and strangers if they are not. Each person picks a position to stand on, yielding a starting configuration. From here, at any point in time, two friends standing on adjacent positions may switch places: we

^{*}This research was initiated at the University of Minnesota Duluth REU and was supported, in part, by NSF-DMS grant 1949884 and NSA Grant H98230-20-1-0009.

call this operation a friendly swap. From the initial configuration, say the n people have a final configuration in mind, and they know it can be reached from the initial configuration by some sequence of friendly swaps. What is the worst-case (over pairs of starting and final configurations) number of friendly swaps that is necessary in order for the n people to achieve the final configuration from the starting configuration?

We may formalize the problem using the following definition.

Definition 1.1 ([DK21]). Let X and Y be simple graphs on n vertices. The *friends-and-strangers graph* of X and Y , denoted $\text{FS}(X, Y)$, is a graph with vertices consisting of all bijections from $V(X)$ to $V(Y)$, with bijections $\sigma, \tau \in \text{FS}(X, Y)$ adjacent if and only if there exists an edge $\{a, b\}$ in X such that

1. $\{\sigma(a), \sigma(b)\} \in E(Y)$,
2. $\sigma(a) = \tau(b)$, $\sigma(b) = \tau(a)$,
3. $\sigma(c) = \tau(c)$ for all $c \in V(X) \setminus \{a, b\}$.

In other words, σ and τ differ precisely on two adjacent vertices of X whose images under σ (and τ) are adjacent in Y . For any such bijections σ, τ , we say that τ is achieved from σ by an (X, Y) -friendly swap.




Figure 1.1: A sequence of (X, Y) -friendly swaps in $\text{FS}(X, Y)$ for the 5-vertex graphs X and Y . Configurations in the bottom row correspond to vertices in $V(\text{FS}(X, Y))$. Two consecutive configurations differ by an (X, Y) -friendly swap, so the corresponding vertices are adjacent in $\text{FS}(X, Y)$.

See Figure 1.1 for an illustration of Definition 1.1 on five-vertex graphs. Defant and Kravitz [DK21], in addition to introducing the framework of friends-and-strangers graphs, described the connected components of $\text{FS}(\text{Path}_n, Y)$ and $\text{FS}(\text{Cycle}_n, Y)$ in terms of the acyclic orientations of \bar{Y} (the complement of Y), and determined both necessary conditions and sufficient conditions for $\text{FS}(X, Y)$ to be connected. In a different paper [Jeo22], we extend their results: [DK21, Corollary 4.14] states that $\text{FS}(\text{Cycle}_n, Y)$ is connected if and only if \bar{Y} is a forest with trees of jointly coprime sizes, and we establish that if X is biconnected (i.e., connected and with no cut vertex) and Y is a graph for which $\text{FS}(\text{Cycle}_n, Y)$ is connected, then $\text{FS}(X, Y)$ is connected, settling [DK21, Conjecture 7.1]. In [Jeo22], we also initiate the study of the girth of friends-and-strangers graphs. Motivated by [KMS84] and connections to molecular programming as seen in [BGYW19], the framework of friends-and-strangers was later generalized by [Mil24] to permit for multiplicities onto vertices, in which many of the main results of [DK21, Wil74] were also generalized accordingly.

A central objective in the study of friends-and-strangers graphs is to determine necessary and sufficient conditions for their connectivity. Indeed, $\text{FS}(X, Y)$ being connected corresponds exactly to the property that one can go between any two configurations in $\text{FS}(X, Y)$ via some sequence of (X, Y) -friendly swaps. Of course, the conditions one may derive will depend upon the assumptions on X and Y under which one works. If one elects to proceed under a regime in which $\text{FS}(X, Y)$ cannot be connected (such as when X and Y are both bipartite; see the discussion around [DK21, Proposition 2.7] and [ADK23, Subsection 2.3] for a parity obstruction which demonstrates why this is the case), one may instead study how small the number of connected components may be under this regime, and the natural question here is to ask for further conditions on X and Y ensuring that $\text{FS}(X, Y)$ achieves the smallest possible number of connected components. As pursued in [DK21, Sections 3 and 4] for (respectively) paths and cycles, one direction of inquiry is to fix (without loss of generality, as we will see in Proposition 2.3(1)) X to be some particular graph, and study the structure of $\text{FS}(X, Y)$ for arbitrary Y : see [DDLW24, Lee22, WC23, Wil74, Zhu24]. It is also very natural to ask extremal and probabilistic questions concerning the connectivity of friends-and-strangers graphs, such as minimum degree conditions on X and Y which ensure that $\text{FS}(X, Y)$ is connected or for threshold probabilities on Erdős-Rényi random graphs X, Y regarding the connectivity of $\text{FS}(X, Y)$: see [ADK23, Ban22, Jeo23, Mil24, WLC23].

The setup proposed by Definition 1.1 is quite general. Indeed, friends-and-strangers graphs serve both as a common natural generalization of many classical combinatorial objects and as a framework which embodies many important problems in discrete mathematics and theoretical computer science. We illustrate this claim with a non-exhaustive listing of relevant examples. The graph $\text{FS}(X, K_n)$ is the Cayley graph of the symmetric group on the vertex set of X generated by the transpositions corresponding to the edges of X ; we refer the reader to [DK21] and the references therein for a comprehensive discussion regarding the relevance of friends-and-strangers graphs within algebraic combinatorics. Letting X be the 4-by-4 grid and Y a star graph, studying $\text{FS}(X, Y)$ is equivalent to studying the configurations and moves that can be performed on the famous 15-puzzle (with the central vertex of the star graph corresponding to the empty tile); see [BK23, DR18, Par15, Yan11] for similar inquiries of a recreational flavor. The works [Naa00, Sta12] both study the structure of the graph $\text{FS}(\text{Path}_n, Y)$ under certain re-

restrictions on Y , while the works [BR99, Rei98] utilize $\text{FS}(\text{Path}_n, Y)$ to investigate the acyclic orientations of \bar{Y} . Asking if X and Y pack [BE78, BJS17, KO09, SS78, Yap88, Yus07] in the graph packing literature is equivalent to asking if there exists an isolated vertex in $\text{FS}(X, Y)$. Studying the token swapping problem [ADK⁺22, BJL⁺23, BMR18, MNO⁺16, YDI⁺15] on the graph X is equivalent to studying distances between configurations in $\text{FS}(X, K_n)$. Finally, as we will briefly touch upon in Subsection 4.4, the interchange process on the graph X [AF02, AK13, Ang03, BD06, CLR10, ES23, Ham15, HS21, Sch05] can be phrased in terms of random walks on $\text{FS}(X, K_n)$.

1.2. Main Results and Organization

Unlike the existing body of work that studies the connectivity of friends-and-strangers graphs, the present paper initiates the study of their diameters, corresponding to the length of the “longest shortest path,” with lengths of shortest paths evaluated over all pairs of vertices. Indeed, the diameter of a connected component of $\text{FS}(X, Y)$ corresponds to the largest number of (X, Y) -friendly swaps necessary to achieve one configuration in the component from another. In a more recreational tone, if we think of $\text{FS}(X, Y)$ as a generalized 15-puzzle, we are asking for the longest solution length for any solvable puzzle involving “board X and rules Y .” The works [ADK23, DK21] both posed the following extremal question, which asks whether the distance between any two configurations in $\text{FS}(X, Y)$ is polynomial in the size of X and Y .

Question 1.2 ([ADK23, DK21]). Does there exist an absolute constant $C > 0$ such that for all n -vertex graphs X and Y , every connected component of $\text{FS}(X, Y)$ has diameter at most n^C ?

In Section 2, we introduce some background that we shall need later in the work. Before tackling the more global Question 1.2, in Section 3, we fix (without loss of generality) X to be a complete, path, or cycle graph, and derive upper bounds on the maximum diameter of a component of $\text{FS}(X, Y)$ in each setting. Our results on paths and cycles address an open problem posed in [DK21, Subsection 7.3]. Furthermore, the discussion therein suggests that one must restrict their attention to rather contrived choices of graphs X and Y in order for $\text{FS}(X, Y)$ to have a component with diameter that is superpolynomial in the size of X and Y , suggesting that Question 1.2 may be challenging to settle via constructive means if it holds in the negative.

In Section 4, we establish the main result of this article, Theorem 1.3, which answers Question 1.2 in the negative. We prove this theorem by constructing, for all integers $L \geq 1$, graphs X_L and Y_L on the same number of vertices: see Figure 1.2 for a schematic diagram of the construction for $L = 3$. The construction is such that the number of vertices of X_L and Y_L is $\Theta(L)$, and there exist two configurations $\sigma_s, \sigma_f \in V(\text{FS}(X_L, Y_L))$ which lie in the same connected component \mathcal{C} of $\text{FS}(X_L, Y_L)$ and for which the distance between σ_s and σ_f is $e^{\Omega(n)}$.

Theorem 1.3. *For all $n \geq 2$, there exist n -vertex graphs X and Y such that $\text{FS}(X, Y)$ has a connected component with diameter $e^{\Omega(n)}$.*

We conclude the work with Section 5, which suggests several open problems and directions for future research.




Figure 1.2: The graphs X_3 and Y_3 .

1.3. Notation

In this article, unless stated otherwise, we assume that all graphs are simple. We employ standard asymptotic notation in this paper. Unless stated otherwise, all asymptotic notation in this paper will be with respect to n . Purely for the sake of completeness, we state the following standard notation.

- The vertex and edge sets of a graph G are denoted by $V(G)$ and $E(G)$, respectively.
- The complement of the graph G is denoted \overline{G} .
- The statement that G and H are isomorphic is written as $G \cong H$.
- For a subset $S \subset V(G)$, we let $G|_S$ denote the induced subgraph of G with vertex set S .
- The open neighborhood of $v \in V(G)$, which is the collection of all neighbors of v , is denoted by $N_G(v)$. The closed neighborhood of $v \in V(G)$ is denoted $N_G[v] = N_G(v) \cup \{v\}$. For a subset of vertices $S \subseteq V(G)$, we let

$$N_G(S) = \bigcup_{v \in S} N_G(v), \quad N_G[S] = \bigcup_{v \in S} N_G[v].$$

- The disjoint of a collection of graphs $\{G_i\}_{i \in I}$, notated $\bigoplus_{i \in I} G_i$, is the graph with vertex set $\bigsqcup_{i \in I} V(G_i)$ and edge set $\bigsqcup_{i \in I} E(G_i)$. This readily extends to expressing a graph as the disjoint union of its connected components.

- The distance $d(v, w)$ between $v, w \in V(G)$ is the length of the shortest path from v to w . The diameter of a component \mathcal{C} of G is $\max_{v, w \in V(\mathcal{C})} d(v, w)$.

Graphs with vertex set $[n] := \{1, \dots, n\}$ that will be relevant later are

- the complete graph K_n , with $E(K_n) := \{\{i, j\} : \{i, j\} \in [n], i \neq j\}$;
- the complete bipartite graph $K_{i,j}$, with $E(K_{i,j}) := \{\{v_1, v_2\} : v_1 \in [i], v_2 \in \{i+1, \dots, i+j\}\}$, which naturally partitions $V(K_{i,j})$ into two sets (henceforth called partite sets);
- the path graph Path_n , with $E(\text{Path}_n) := \{\{i, i+1\} : i \in [n-1]\}$;
- the cycle graph Cycle_n , with $E(\text{Cycle}_n) := \{\{i, i+1\} : i \in [n-1]\} \cup \{\{n, 1\}\}$;
- the star graph $\text{Star}_n := K_{1, n-1}$.

2. Background

In this section, we introduce some background and summarize results from prior work that will be relevant later in the present paper, particularly in Section 3. Throughout this section and Section 3, we will assume that the vertex set of all graphs is $[n]$, with edge sets as in Subsection 1.3. Note that if both $V(X)$ and $V(Y)$ are $[n]$, then the vertices of $\text{FS}(X, Y)$ are the elements of \mathfrak{S}_n , the symmetric group of degree n .

2.1. Acyclic Orientations

An *orientation* of a graph G is an assignment of a direction to every edge of G , and an *acyclic orientation* of G is an orientation with no directed cycles. Denote the set of all acyclic orientations of G by $\text{Acyc}(G)$. We will be interested in operations on acyclic orientations of G called *flips* and *double-flips*, as defined in [DK21]. Notably, it was shown in [DK21, Theorem 4.7] that double-flips on acyclic orientations in $\text{Acyc}(\overline{Y})$ are paramount in describing the connected components of $\text{FS}(\text{Cycle}_n, Y)$.

Letting $\alpha \in \text{Acyc}(G)$, converting a source of α into a sink or a sink of α into a source by reversing the directions of all its incident edges results in another acyclic orientation α' of G . We call such an operation a *flip*, and we say that α and α' are *flip equivalent*, denoted $\alpha \sim \alpha'$. In the literature, the equivalence classes in $\text{Acyc}(G)/\sim$ are called *toric acyclic orientations*; we refer the interested reader to [Che10, DMR16, MM11, Pre86, Spe09] for related reading. We will further say that we perform an *inflip* on α if we convert a source into a sink (the direction of all incident edges “go into” the new sink), and an *outflip* if we convert a sink into a source.¹

Similarly, flipping a nonadjacent source and sink of α into (respectively) a sink and a source results in another acyclic orientation α'' of G : we call such an operation a *double-flip*, and we say α and α'' are *double-flip equivalent*, denoted $\alpha \approx \alpha''$. It is easy to show that \sim and \approx are equivalence relations on $\text{Acyc}(G)$. We denote the set of double-flip equivalence classes

¹In particular, we may apply these operations to isolated vertices.

of $\text{Acyc}(G)$ by $\text{Acyc}(G)/\approx$, and denote the *double-flip equivalence class* for which α is a representative by $[\alpha]_{\approx}$.

Assume $V(G) = [n]$, and take $\alpha \in \text{Acyc}(G)$. Associated to the acyclic orientation α is a poset $P_{\alpha} = ([n], \leq_{\alpha})$, where $i \leq_{\alpha} j$ if and only if there exists a directed path from i to j in α . We define a *linear extension* of P_{α} to be any permutation $\sigma \in \mathfrak{S}_n$ such that $\sigma^{-1}(i) \leq \sigma^{-1}(j)$ whenever $i \leq_{\alpha} j$. We let $\mathcal{L}(\alpha)$ denote the collection of linear extensions of P_{α} . For any $\sigma \in \mathfrak{S}_n$, it is not hard to see that there exists a unique acyclic orientation $\alpha_G(\sigma) \in \text{Acyc}(G)$ for which $\sigma \in \mathcal{L}(\alpha_G(\sigma))$, and that this acyclic orientation is the result of directing each edge $\{i, j\} \in E(G)$ from i to j if and only if $\sigma^{-1}(i) < \sigma^{-1}(j)$. It is also not hard to see that the poset P_{α} associated to $\alpha \in \text{Acyc}(G)$ has a linear extension (e.g., for $i \in [n]$, we can construct a linear extension σ by setting $\sigma^{-1}(i)$ to be a source of α , then removing the source and all incident edges from α ; in an abuse of notation,² we understand α here as being mutated over the course of this greedy algorithm). We write

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

We refer the reader to [DK21, Section 4] for a more comprehensive discussion regarding why these notions are of importance in the study of friends-and-strangers graphs (though this is illuminated in passing in Subsection 2.2 and in the arguments of Section 3).

For a graph G and acyclic orientation $\alpha \in \text{Acyc}(G)$, we can partition the directed edges of any cycle subgraph \mathcal{C} of G into \mathcal{C}_{α}^{-} and \mathcal{C}_{α}^{+} , corresponding to edges directed in one of two possible directions under α in \mathcal{C} . The article [Pre86] studied precisely when an acyclic orientation could be reached from another by a sequence of inflips or outflips, while [Pro21] extends this result by providing an upper bound on the number of inflips or outflips necessary to reach α from α' whenever $\alpha \sim \alpha'$.

Lemma 2.1 ([Pre86, Pro21]). *For $\alpha, \alpha' \in \text{Acyc}(G)$, α' can be reached from α by a sequence of inflips if and only if for every cycle subgraph \mathcal{C} of G , $|\mathcal{C}_{\alpha}^{-}| = |\mathcal{C}_{\alpha'}^{-}|$. Furthermore, whenever this is the case, α' can be reached from α by a sequence of at most $\binom{n}{2}$ inflips. Similarly, α' can be reached from α by a sequence of outflips if and only if for every cycle subgraph \mathcal{C} of G , $|\mathcal{C}_{\alpha}^{+}| = |\mathcal{C}_{\alpha'}^{+}|$. Furthermore, whenever this is the case, α' can be reached from α by a sequence of at most $\binom{n}{2}$ outflips.*

We build on Lemma 2.1. The following proposition establishes that we could have defined flip equivalence strictly with respect to inflips or outflips, as this would have resulted in the same notion.

Proposition 2.2. *Acyclic orientations $\alpha, \alpha' \in \text{Acyc}(G)$ are flip equivalent if and only if α' can be reached from α by a sequence of inflips. Similarly, $\alpha \sim \alpha'$ if and only if α' can be reached from α by a sequence of outflips.*

Proof. The statement that α' is reachable from α via a sequence of inflips (or outflips) implying $\alpha \sim \alpha'$ is immediate. To prove the converse, notice that for any cycle subgraph \mathcal{C} of G and

²We will commit similar abuses of notation in Section 3. They should not raise any confusion when invoked.

acyclic orientations $\alpha, \alpha' \in \text{Acyc}(G)$ for which α' can be reached from α by a flip, $|\mathcal{C}_\alpha^-| = |\mathcal{C}_{\alpha'}^-|$. Thus, if $\alpha \sim \alpha'$, then $|\mathcal{C}_\alpha^-| = |\mathcal{C}_{\alpha'}^-|$, so α' can be reached from α via a sequence of inflips (or outflips). \square

2.2. Background on Friends-and-Strangers Graphs

We mention those general properties of friends-and-strangers graphs that we will need later in the article. We refer the reader to [DK21, Section 2] for a thorough treatment of the general properties of friends-and-strangers graphs.

Proposition 2.3 ([DK21, Proposition 2.6]). *The following properties hold.*

1. *Definition 1.1 is symmetric with respect to X and Y : we have that $\text{FS}(X, Y) \cong \text{FS}(Y, X)$.*
2. *The graph $\text{FS}(X, Y)$ is bipartite.*
3. *If X or Y is disconnected, or if X and Y are connected graphs on $n \geq 3$ vertices and each have a cut vertex, then $\text{FS}(X, Y)$ is disconnected.*

The definitions concerning acyclic orientations that were introduced in Subsection 2.1 were observed in [DK21] to be central in describing the structure of the connected components of $\text{FS}(\text{Path}_n, Y)$ and $\text{FS}(\text{Cycle}_n, Y)$, which are the graphs we will be interested in during Section 3. Specifically, we have the following theorems.

Theorem 2.4 ([DK21, Theorem 3.1]). *Let $\alpha \in \text{Acyc}(\overline{Y})$. Take any linear extension $\sigma \in \mathcal{L}(\alpha)$, and let H_α denote the connected component of $\text{FS}(\text{Path}_n, Y)$ which contains σ . Then*

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

and $V(H_\alpha) = \mathcal{L}(\alpha)$. In particular, H_α is independent of the choice of σ .

Theorem 2.5 ([DK21, Theorem 4.7]). *Let $\alpha \in \text{Acyc}(\overline{Y})$. Take any linear extension $\sigma \in \mathcal{L}([\alpha]_{\approx})$, and let $H_{[\alpha]_{\approx}}$ denote the connected component of $\text{FS}(\text{Cycle}_n, Y)$ which contains σ . Then*

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

and $V(H_{[\alpha]_{\approx}}) = \mathcal{L}([\alpha]_{\approx})$. In particular, $H_{[\alpha]_{\approx}}$ is independent of the choice of σ .

Defant and Kravitz [DK21] also determined precisely when $\text{FS}(\text{Cycle}_n, Y)$ is connected. The coprimality condition on the sizes of the components of \overline{Y} in Theorem 2.6 may seem surprising at first glance. We refer the reader to the discussion around [DK21, Corollary 4.12] and [DK21, Corollary 4.14] to see where this condition emerges and why it is a natural one.

Theorem 2.6 ([DK21, Corollary 4.14]). *Let Y be a graph on $n \geq 3$ vertices. Then $\text{FS}(\text{Cycle}_n, Y)$ is connected if and only if \overline{Y} is a forest with trees $\mathcal{T}_1, \dots, \mathcal{T}_r$ such that $\gcd(|V(\mathcal{T}_1)|, \dots, |V(\mathcal{T}_r)|) = 1$.*

3. Diameters of $\text{FS}(X, Y)$ with One Graph Fixed

Before investigating (and settling) the more global question of whether or not the diameters of connected components of friends-and-strangers graphs are polynomially bounded (in the sense posed by Question 1.2), we begin by restricting our study by choosing one of the two graphs X and Y to come from a natural family of graphs, and then establish bounds on the diameter of any connected component of $\text{FS}(X, Y)$.

3.1. Complete Graphs

We begin by setting $Y = K_n$. Take any two configurations $\sigma, \tau \in V(\text{FS}(X, K_n))$ that lie in the same connected component. Consider the following iterative algorithm, applied starting from σ and proceeding sequentially on $i \in [n]$. In an abuse of notation, σ is understood to be mutated over the course of this algorithm as we perform (X, K_n) -friendly swaps to modify its mappings.

1. If $\sigma(i) = \tau(i)$, do nothing.
2. If $\sigma(i) \neq \tau(i)$, swap $\tau(i)$ onto i along a simple path, then swap $\sigma(i)$ back along the simple path that $\tau(i)$ traversed.

It is straightforward to prove via induction that at the beginning of any iteration $i \in [n]$, $\sigma(i)$ and $\tau(i)$ lie upon the same connected component of X (so that the algorithm may always proceed), and that $\sigma(j) = \tau(j)$ for all $j < i$. Thus, $\sigma = \tau$ when the algorithm terminates after $n - 1$ iterations (it must be that $\sigma(n) = \tau(n)$ at the beginning of the n^{th} iteration). For any iteration $i \in [n]$, step (2) requires at most $n - 1$ (X, K_n) -friendly swaps to move $\tau(i)$ onto i , and at most $n - 2$ (X, K_n) -friendly swaps to move $\sigma(i)$ back. This establishes that the diameter of any component of $\text{FS}(X, K_n)$ is therefore at most $(n-1)((n-1)+(n-2)) = 2n^2 - 5n + 3 = O(n^2)$.

Finding the exact distance between two configurations in $\text{FS}(X, K_n)$ is known as the *token swapping problem on X* in the theoretical computer science literature. The $O(n^2)$ bound on the diameter of any component of $\text{FS}(X, K_n)$ is well known, and we also have a bound of $\Omega(n^2)$ on the diameter of any component of $\text{FS}(X, K_n)$ for particular choices of X (e.g., see Remark 3.4). In general, computing exact distances between two configurations in $\text{FS}(X, K_n)$, as well as the diameters of its connected components, is challenging, even when imposing additional assumptions on X (e.g., see [BJL⁺23, YDI⁺15]). There do exist, however, exact polynomial-time algorithms which solve the token swapping problem for a number of choices of X , including cliques [Cay49], paths [Jer85], stars [PV90], cycles [KSY19, vBSY16], and complete bipartite graphs [YDI⁺15]. See Subsection 5.5 for additional discussion regarding matters of hardness and approximation.

3.2. Paths

In this subsection, we fix $X = \text{Path}_n$. We begin by introducing a notion which will serve as a monovariant in the proof of Proposition 3.2.

Definition 3.1. For $\sigma, \tau \in \mathfrak{S}_n$, call the ordered pair (i, j) ($i, j \in [n], i < j$) a (σ, τ) -inversion if either

1. $\sigma^{-1}(i) < \sigma^{-1}(j)$ and $\tau^{-1}(j) < \tau^{-1}(i)$,
2. $\sigma^{-1}(j) < \sigma^{-1}(i)$ and $\tau^{-1}(i) < \tau^{-1}(j)$.

Denote the number of (σ, τ) -inversions by $\text{inv}(\sigma, \tau)$.

In other words, the ordered pair (i, j) is a (σ, τ) -inversion if the relative ordering of the inverse images of i, j under σ is opposite that of τ . If (without loss of generality) τ is the identity permutation, then $\text{inv}(\sigma, \tau) = \text{inv}(\sigma)$, the number of inversions of σ . It also follows immediately that $\text{inv}(\sigma, \tau) = 0$ if and only if $\sigma = \tau$.

Proposition 3.2. Take $\alpha \in \text{Acyc}(\overline{Y})$, and let H_α denote the corresponding connected component of $\text{FS}(\text{Path}_n, Y)$. Let $P_\alpha = ([n], \leq_\alpha)$ be the poset on $[n]$ for which $i \leq_\alpha j$ if and only if there exists a directed path from i to j in \overline{Y} under α . Then $\text{diam}(H_\alpha) \leq \binom{n}{2} - p_\alpha$, where p_α denotes the number of comparable ordered pairs (i, j) with $i, j \in [n], i < j$ in P_α .

Proof. We will show for any $\sigma, \tau \in V(H_\alpha)$ that $d(\sigma, \tau) = \text{inv}(\sigma, \tau)$. Any (Path_n, Y) -friendly swap reduces the number of (σ, τ) -inversions by at most one, so $d(\sigma, \tau) \geq \text{inv}(\sigma, \tau)$. Now consider the following variant of the bubble sort algorithm, which we perform beginning from $\sigma = \sigma(1)\sigma(2) \cdots \sigma(n)$. Say $\sigma(i_1) = \tau(1)$, and swap $\sigma(i_1)$ down to position 1, yielding σ_1 with $\sigma_1(1) = \tau(1)$. Now, say $\sigma(i_2) = \tau(2)$ (with $i_2 \geq 2$), and swap $\sigma(i_2)$ down to position 2, yielding σ_2 with $\sigma_2(j) = \tau(j)$ for $j \in [2]$; continue until we achieve $\sigma_n = \tau$. It is immediate that the execution of any swap performed during this algorithm would decrement $\text{inv}(\sigma, \tau)$ by 1. Furthermore, any proposed swap in this algorithm can be executed, i.e., involves two elements which comprise an edge in Y . Indeed, Theorem 2.4 yields $\sigma, \tau \in V(H_\alpha) = \mathcal{L}(\alpha)$, but the existence of a swap in this algorithm that cannot be executed would yield $\alpha_{\overline{Y}}(\sigma) \neq \alpha_{\overline{Y}}(\tau)$ (if the proposed swap fails to be an edge in Y , it is an edge in \overline{Y} , and would be directed in opposite directions under $\alpha_{\overline{Y}}(\sigma)$ and $\alpha_{\overline{Y}}(\tau)$ because the two elements comprising the swap constitute a (σ, τ) -inversion), which is a contradiction. Thus, $d(\sigma, \tau) = \text{inv}(\sigma, \tau)$. If $(i, j) \in P_\alpha$, it follows from $\sigma, \tau \in \mathcal{L}(\alpha)$ that $\sigma^{-1}(i) < \sigma^{-1}(j)$ and $\tau^{-1}(i) < \tau^{-1}(j)$, so (i, j) is not a (σ, τ) -inversion. Thus, $d(\sigma, \tau) = \text{inv}(\sigma, \tau) \leq \binom{n}{2} - p_\alpha$, and therefore $\text{diam}(H_\alpha) \leq \binom{n}{2} - p_\alpha$. \square

Certainly, the two vertices incident to an edge of \overline{Y} are comparable in the poset $P_\alpha = ([n], \leq_\alpha)$ for any $\alpha \in \text{Acyc}(\overline{Y})$. This yields the following statement, as $\binom{n}{2} - p_\alpha \leq \binom{n}{2} - |E(\overline{Y})| = |E(Y)|$. For simplicity, we appeal to Theorem 3.3, rather than Proposition 3.2, in forthcoming arguments.

Theorem 3.3. The diameter of any connected component of $\text{FS}(\text{Path}_n, Y)$ is at most $|E(Y)|$.

Remark 3.4. It is not hard to see that $\text{FS}(\text{Path}_n, K_n)$ is connected (e.g., for any two configurations $\sigma, \tau \in V(\text{FS}(\text{Path}_n, K_n))$, the algorithm from Subsection 3.1 yields a path between σ and τ). From the proof of Proposition 3.2, we have for any $\sigma, \tau \in V(\text{FS}(\text{Path}_n, K_n))$ that $d(\sigma, \tau) = \text{inv}(\sigma, \tau) \leq \binom{n}{2}$, and $\text{inv}(\sigma, \tau) = \binom{n}{2}$ when τ is the “reverse” of σ (i.e., $\tau(i) = \sigma(n - i + 1)$ for all $i \in [n]$). So $\text{diam}(\text{FS}(\text{Path}_n, K_n)) = \binom{n}{2}$. Combined with

Proposition 3.2, this establishes that the maximum diameter of a component of $\text{FS}(\text{Path}_n, K_n)$ is $\Omega(n^2)$, and thus $\Theta(n^2)$. Thus, there exist families of n -vertex graphs Y for which the maximum diameter of a component of $\text{FS}(\text{Path}_n, Y)$ has diameter $\Theta(n^2)$. The same can be said for $\text{FS}(K_n, Y)$.

Remark 3.5. The upper bound of $\binom{n}{2} - p_\alpha$ on $\text{diam}(H_\alpha)$ in Theorem 3.2 corresponds to the number of ordered pairs (i, j) ($i < j; i, j \in [n]$) that are incomparable in the poset $P_\alpha = ([n], \leq_\alpha)$. It follows from the proof of Proposition 3.2 that for arbitrary $\sigma, \tau \in V(H_\alpha) = \mathcal{L}(\alpha)$, any (σ, τ) -inversion must be a pair of incomparable elements in P_α , and $d(\sigma, \tau) = \text{inv}(\sigma, \tau)$. We now apply these observations to show that the upper bound on $\text{diam}(H_\alpha)$ fails to be sharp: Figure 3.1 provides an illustration of our construction. For $n = 6$, consider the graph shown in Figure 3.1a, whose complement is shown in Figure 3.1b. We will take $\alpha \in \text{Acyc}(\bar{Y})$ to be an acyclic orientation for which the edges in this connected component are oriented as in Figure 3.1c.




Figure 3.1: The construction we used to show that the bound given in Proposition 3.2 fails to be sharp in general.

Assume (towards a contradiction) that there exist $\sigma, \tau \in V(H_\alpha) = \mathcal{L}(\alpha)$ for which $d(\sigma, \tau) = \text{inv}(\sigma, \tau) = \binom{n}{2} - p_\alpha$, so that all pairs of incomparable elements in P_α are (σ, τ) -inversions. Any two elements in $\{1, 2, 3\}$ are incomparable in P_α , so the relative ordering of $\{1, 2, 3\}$ in σ must be the relative ordering of $\{1, 2, 3\}$ in τ reversed. Without loss of generality, assume σ has relative ordering $1 \rightarrow 2 \rightarrow 3$, so τ has $3 \rightarrow 2 \rightarrow 1$. Since $\sigma, \tau \in \mathcal{L}(\alpha)$, the element 4 follows vertex 2 in both σ and τ , so $(2, 4)$ is not a (σ, τ) -inversion. But $(2, 4)$ is incomparable in P_α , a contradiction.

3.3. Cycles

In this subsection, we fix $X = \text{Cycle}_n$. The setting $Y = K_n$ has been studied in the context of circular permutations [Kim16, vBSY16]. In particular, [Kim16, Procedure 3.6] provides an algorithm that achieves the minimal number of (Cycle_n, K_n) -friendly swaps between any two permutations in \mathfrak{S}_n . Extracting these results yields that the diameter of $\text{FS}(\text{Cycle}_n, K_n)$ is $\lfloor n^2/4 \rfloor$. In the spirit of Remark 3.4, it follows that there exist families of n -vertex graphs Y for which $\text{FS}(\text{Cycle}_n, Y)$ has diameter $\Theta(n^2)$, and it is worth asking what conditions on Y yield that the maximum diameter of a connected component of $\text{FS}(\text{Cycle}_n, Y)$ is at most quadratic in n . In this direction, we have the following proposition.

Proposition 3.6. *If Y has an isolated vertex or $|E(Y)| \leq n - 2$, then the diameter of any connected component of $\text{FS}(\text{Cycle}_n, Y)$ is at most $|E(Y)|$.*

Proof. Consider any $\sigma, \tau \in V(\text{FS}(\text{Cycle}_n, Y))$ which lie in the same component. If Y has an isolated vertex v , then it must be that $\sigma^{-1}(v)$ remains fixed over any sequence of (Cycle_n, Y) -friendly swaps from σ to τ . Thus, it must be that any path from σ to τ in $\text{FS}(\text{Cycle}_n, Y)$ is a path in

$$\text{FS}(\text{Cycle}_n |_{V(\text{Cycle}_n) \setminus \{\sigma^{-1}(v)\}}, Y_{V(Y) \setminus \{v\}}),$$

from which the result follows from Theorem 3.3. For the setting $|E(Y)| \leq n - 2$, we will show that any $\sigma, \tau \in V(\text{FS}(\text{Cycle}_n, Y))$ in the same connected component will remain in the same component after removing some edge from Cycle_n , from which the desired result again follows immediately from Theorem 3.3. Assume (towards a contradiction) that every path from σ to τ in $\text{FS}(\text{Cycle}_n, Y)$ involves a swap over every edge in $E(\text{Cycle}_n)$. Consider a shortest path $\Sigma = \{\sigma_i\}_{i=0}^\lambda$ from σ or τ , which has that $\sigma_0 = \sigma$ and $\sigma_\lambda = \tau$, and $\lambda \geq n$ by the assumption. Consider the subsequence $\{\sigma_i\}_{i=0}^{n-1}$ consisting of the first $n - 1$ (Cycle_n, Y) -friendly swaps of Σ . This must be a shortest path from σ to σ_{n-1} in $\text{FS}(\text{Cycle}_n, Y)$, and swaps upon at most $n - 1$ edges of Cycle_n : say $e \in E(\text{Cycle}_n)$ is an edge upon which a swap does not occur, and let Cycle_n^{-e} be Cycle_n with this edge e removed. Then $\{\sigma_i\}_{i=0}^{n-1}$ is a shortest path from σ to σ_{n-1} in $\text{FS}(\text{Cycle}_n^{-e}, Y)$ with length $n - 1$. This contradicts Theorem 3.3, which yields $d(\sigma, \sigma_{n-1}) \leq |E(Y)| \leq n - 2$. \square

We were unable to extend the $O(n^2)$ bound from Proposition 3.6 to general Y , although we suspect that this is the truth (see Subsection 5.2). However, the existence of a universal constant $C > 0$ such that the maximum diameter of a component of $\text{FS}(\text{Cycle}_n, Y)$ is $O(n^C)$ remains highly desirable. In conjunction with Theorem 2.6, the following theorem yields such a result whenever $\text{FS}(\text{Cycle}_n, Y)$ is connected.

Theorem 3.7. *Let Y be a graph on $n \geq 3$ vertices, and let n_1, \dots, n_r denote the sizes of the components of \bar{Y} . If $\gcd(n_1, \dots, n_r) = 1$, then any component of $\text{FS}(\text{Cycle}_n, Y)$ has diameter at most $4n^3 + |E(Y)|$.*

Proof. Certainly, $r \geq 2$. Without loss of generality, we assume that $n_1 \leq \dots \leq n_r$, and we denote the corresponding components of \bar{Y} by $\bar{Y}_1, \dots, \bar{Y}_r$, respectively. For $\alpha \in \text{Acyc}(\bar{Y})$, we let α_i denote the acyclic orientation induced by α on \bar{Y}_i . We now fix $\alpha, \alpha'' \in \text{Acyc}(\bar{Y})$ such that $\alpha \approx \alpha''$. Before studying distances in $\text{FS}(\text{Cycle}_n, Y)$, we will first bound the number of double-flips necessary to reach α'' from α . Certainly, $\alpha_i \sim \alpha''_i$ for all $i \in [r]$, and by Proposition 2.2, we can reach α''_i in no more than $\binom{n_i}{2}$ inflips or outflips from α_i . Observe that for any α''_i , we may return to α''_i by applying a different sequence of n_i inflips; see Figure 3.2 for an illustration. Indeed, take a linear extension $\sigma \in \mathcal{L}(\alpha''_i)$, labeled $\sigma = \sigma(1)\sigma(2) \cdots \sigma(n_i)$, and perform an inflip on α''_i by converting the source $\sigma(1)$ into a sink, so that $\sigma(2) \cdots \sigma(n_i)\sigma(1)$ is a linear extension of the poset associated to the resulting acyclic orientation in $\text{Acyc}(\bar{Y}_i)$. Performing n_i inflips on α''_i in this manner returns σ as a linear extension of the poset associated to the resulting acyclic orientation: since there exists a unique acyclic orientation $\alpha_{\bar{Y}}(\sigma) \in \text{Acyc}(\bar{Y}_i)$




Figure 3.2: An example of a sequence of n inflips which takes an acyclic orientation α of an n -vertex graph back to itself. We demonstrate on a 4-vertex graph. The permutations $\begin{pmatrix} 1 & 2 & 3 & 4 \\ 4 & 1 & 2 & 3 \end{pmatrix}$, $\begin{pmatrix} 1 & 2 & 3 & 4 \\ 1 & 2 & 3 & 4 \end{pmatrix}$, $\begin{pmatrix} 1 & 2 & 3 & 4 \\ 2 & 3 & 4 & 1 \end{pmatrix}$, $\begin{pmatrix} 1 & 2 & 3 & 4 \\ 3 & 4 & 1 & 2 \end{pmatrix}$ are linear extensions of the posets associated with the first four acyclic orientations shown, respectively. The first and fifth acyclic orientations are the same.

for which $\sigma \in \mathcal{L}(\alpha_{\overline{r}}(\sigma))$, this acyclic orientation must be α'_i . Similarly, we can return to α''_i by applying a sequence of n_i outflips.

Recalling that a double-flip applied to an acyclic orientation involves flipping a nonadjacent source and sink into (respectively) a sink and source, we thus proceed as follows. Starting from the acyclic orientation α , perform a sequence of double-flips that act as inflips on sources in α_r and outflips on sinks in $\alpha_1, \dots, \alpha_{r-1}$ until we have reached $\alpha''_1, \dots, \alpha''_r$ at least once. Specifically, begin by performing inflips on α_r and outflips on α_1 until we either reach α''_1 (at which point we begin performing outflips on sinks in α_2) or α''_r (at which point we begin performing inflips on sources in α_r as described previously to return to α_r every n_r inflips). If we reach $\alpha''_1, \dots, \alpha''_{r-1}$ prior to α''_r , then perform outflips on sinks in α''_1 (returning to α''_1 every n_1 outflips) until α''_r is reached: from here, pair these outflips on sinks with inflips on sources in α''_r until we retain α''_1 . Otherwise, we reach α''_r prior to $\alpha''_1, \dots, \alpha''_{r-1}$, for which α''_r will be “offset” once we have $\alpha''_1, \dots, \alpha''_{r-1}$, since we are performing inflips on sources which return to α''_r every n_r inflips. In either case, call the resulting acyclic orientation $\tilde{\alpha}$, which satisfies $\tilde{\alpha}_i = \alpha''_i$ for all $i \in [r - 1]$ while $\tilde{\alpha}_r$ differs from α''_r by some offset $0 \leq c < n_r$. By tracing the preceding description and recalling Proposition 2.2, it follows that the number of double-flips we perform to reach $\tilde{\alpha}$ from α is bounded above by

$$\max \left\{ \binom{n_r}{2} + n_1, \sum_{i=1}^{r-1} \binom{n_i}{2} \right\} \leq \sum_{i=1}^r n_i^2 \leq \left(\sum_{i=1}^r n_i \right)^2 = n^2.$$

By Bézout’s Lemma (recall that $\gcd(n_1, \dots, n_r) = 1$), there exist integers $0 \leq d_1, \dots, d_{r-1} < n_r$ such that


$$d_1 n_1 + \dots + d_{r-1} n_{r-1} \equiv n_r - c \pmod{n_r}.$$

Thus, from $\tilde{\alpha}$, we can reach α'' by performing $d_i n_i$ outflips on $\tilde{\alpha}_i = \alpha''_i$ for $i \in [r - 1]$ (returning to $\tilde{\alpha}_i = \alpha''_i$ every n_i outflips), while performing inflips on $\tilde{\alpha}_r$ as discussed to reach α''_r . The number of double-flips we perform to reach α'' from $\tilde{\alpha}$ is therefore bounded above by

$$\sum_{i=1}^{r-1} d_i n_i \leq \max \{d_1, \dots, d_{r-1}\} \left(\sum_{i=1}^{r-1} n_i \right) \leq n_r n \leq n^2,$$

so at most $2n^2$ double-flips are necessary to reach α'' from α .

We now turn to bounding $d(\sigma, \tau)$ for configurations $\sigma, \tau \in V(\text{FS}(\text{Cycle}_n, Y))$ in the same connected component. By Theorem 2.5, we have that $\sigma, \tau \in \mathcal{L}([\alpha]_{\approx})$ for some $[\alpha]_{\approx} \in \text{Acyc}(\overline{Y})/\approx$. Denote $\alpha = \alpha_{\overline{Y}}(\sigma)$ and $\alpha'' = \alpha_{\overline{Y}}(\tau)$. By the preceding discussion, we can reach α'' from α in $\lambda \leq 2n^2$ double-flips, yielding a sequence of acyclic orientations $\Sigma = \{\alpha_i\}_{i=0}^{\lambda}$ in the equivalence class $[\alpha]_{\approx}$ with $\alpha_0 = \alpha$ and $\alpha_{\lambda} = \alpha''$. From Σ , we will now construct a sequence of (Cycle_n, Y) -friendly swaps which we can apply on σ ; see Figure 3.3 for an illustration. If the double-flip we performed to reach α_1 from α inflips the source v and outflips the sink w in α , it follows from $\sigma \in \mathcal{L}(\alpha)$ that for any $i < \sigma^{-1}(v)$, $\{\sigma(i), v\} \in E(Y)$. Indeed, if we had that $\{\sigma(i), v\} \in E(\overline{Y})$, v being a source in α would imply that this edge is directed from v to $\sigma(i)$ in α , contradicting $\sigma \in \mathcal{L}(\alpha)$. Similarly, for any $j > \sigma^{-1}(w)$, $\{\sigma(j), w\} \in E(Y)$. Thus, we can swap v to 1 and w to n in no more than $2n - 3$ (Cycle_n, Y) -friendly swaps: it is easy to check that the resulting configuration remains in $\mathcal{L}(\alpha)$. Then we perform a (Cycle_n, Y) -friendly swap which swaps v and w along the edge $\{1, n\}$ ($\{v, w\} \notin E(\overline{Y})$ by the definition of a double-flip, so $\{v, w\} \in E(Y)$). It is also straightforward to check that the configuration σ_1 resulting from this interchange is now in $\mathcal{L}(\alpha_1)$.



(a) Acyclic orientations $\alpha, \alpha_1 \in \text{Acyc}(\overline{Y})$.

(b) The corresponding sequence of (Cycle_n, Y) -friendly swaps we construct.

Figure 3.3: The sequence of (Cycle_n, Y) -friendly swaps that we construct corresponding to $\alpha, \alpha_1 \in \text{Acyc}(\overline{Y})$ that are double-flip equivalent. We demonstrate on 5-vertex graphs. The topmost bijection is in $\mathcal{L}(\alpha)$. We inflip $v = 2$ and outflip $w = 5$ to reach α'' from α : swapping v left to 1, then w right to 5, then swapping v and w along $\{1, 5\}$ yields a permutation in $\mathcal{L}(\alpha_1)$.

Proceed similarly through all λ double-flips, and call the resulting configuration $\tilde{\sigma}$: this configuration satisfies $\tilde{\sigma} \in \mathcal{L}(\alpha'')$. Since $\tilde{\sigma}, \tau \in \mathcal{L}(\alpha'')$, it follows from Theorem 2.4 that $\tilde{\sigma}, \tau$ lie in the same component of $\text{FS}(\text{Path}_n, Y)$ (specifically, the copy of Path_n in Cycle_n resulting from excluding the edge $\{1, n\}$). By Theorem 3.3, we can now reach τ from $\tilde{\sigma}$ by performing no more than $|E(Y)|$ (Cycle_n, Y) -friendly swaps. Altogether, we have that

$$d(\sigma, \tau) \leq 2n^2 \cdot 2n + |E(Y)| = 4n^3 + |E(Y)|,$$

so at most $4n^3 + |E(Y)|$ (Cycle_n, Y) -friendly swaps are necessary to reach τ from σ . \square

Corollary 3.8. *For $n \geq 3$, if $\text{FS}(\text{Cycle}_n, Y)$ is connected, $\text{diam}(\text{FS}(\text{Cycle}_n, Y)) \leq 4n^3 + |E(Y)|$.*

Theorem 3.7 can now be invoked to establish the following general bound on the diameter of any connected component of $\text{FS}(\text{Cycle}_n, Y)$, where Y is arbitrary. This proves that, in the sense of Question 1.2, the diameter of $\text{FS}(\text{Cycle}_n, Y)$ is polynomially bounded.

Theorem 3.9. *The diameter of any component of $\text{FS}(\text{Cycle}_n, Y)$ is at most $8n^4(1 + o(1))$.*

Proof. Consider two configurations $\sigma, \tau \in V(\text{FS}(\text{Cycle}_n, Y))$ in the same connected component. We construct an $(n + 1)$ -vertex graph Y' by adding a vertex v to Y that is adjacent to all vertices in $V(Y)$, so Y' has a spanning star subgraph with central vertex v . Define bijections $\sigma', \tau' \in V(\text{FS}(\text{Cycle}_{n+1}, Y'))$ by

$$\sigma'(i) = \begin{cases} \sigma(i) & i \in [n], \\ v & i = n + 1, \end{cases} \quad \tau'(i) = \begin{cases} \tau(i) & i \in [n], \\ v & i = n + 1. \end{cases}$$

The configurations σ', τ' are in the same component of $\text{FS}(\text{Cycle}_{n+1}, Y')$. Indeed, from a sequence of (Cycle_n, Y) -friendly swaps Σ_1 from σ to τ of shortest length, we can construct a sequence Σ'_1 of (Cycle_{n+1}, Y') -friendly swaps from σ' to τ' by replacing every swap in Σ_1 which occurs along $\{1, n\} \in E(\text{Cycle}_n)$ by a sequence of three swaps along the following edges in $E(\text{Cycle}_n)$:

$$\{n, n + 1\}, \{1, n + 1\}, \{n, n + 1\}.$$

It is straightforward to confirm that Σ'_1 is a path from σ' to τ' , constructed from Σ_1 by ‘‘crossing’’ the vertex v as needed. Since Y' has a spanning star subgraph, $\overline{Y'}$ has an isolated vertex, so it follows immediately that the components of $\overline{Y'}$ have jointly coprime size. So by Theorem 3.7,

$$d(\sigma', \tau') \leq 4(n + 1)^3 + |E(Y')| \leq 4(n + 1)^3 + \binom{n + 1}{2} = 4n^3(1 + f(n)),$$

where $f(n) = o(1)$. Let Σ'_2 be a sequence of swaps from σ' to τ' of length at most $4n^3(1 + f(n))$. We construct Σ_2 from Σ'_2 by removing all (Cycle_{n+1}, Y') -friendly swaps involving v : it is straightforward to notice that Σ_2 yields a path of length at most $4n^3(1 + f(n))$ from σ to some cyclic rotation τ_* of τ , i.e., $d(\sigma, \tau_*) \leq 4n^3(1 + f(n))$. Towards a contradiction, assume $d(\tau, \tau_*) \geq 8n^4(1 + 2f(n)) + n$, and let v_1, \dots, v_{n+1} be vertices along a shortest path from τ to τ_* satisfying $d(v_i, v_{i+1}) > 8n^3(1 + 2f(n))$ for all $i \in [n]$. Such vertices v_i exist due to our assumption on $d(\tau, \tau_*)$. By appealing to the same argument as above, we deduce that there exist cyclic rotations $\sigma_1, \dots, \sigma_{n+1}$ of σ such that $d(v_i, \sigma_i) \leq 4n^3(1 + f(n))$ for all $i \in [n + 1]$. Since there exist n distinct rotations of σ , the pigeonhole principle yields the existence of $i \neq j$ for which

$$d(v_i, v_j) \leq d(v_i, \sigma') + d(\sigma', v_j) \leq 4n^3(1 + f(n)) + 4n^3(1 + f(n)) \leq 8n^3(1 + 2f(n))$$

for some rotation σ' of σ , which is a contradiction. Therefore, we conclude that

$$d(\sigma, \tau) \leq d(\sigma, \tau_*) + d(\tau_*, \tau) \leq 4n^3(1 + f(n)) + 8n^3(1 + 2f(n))(n + 1) = 8n^4(1 + o(1)).$$

The desired result now follows immediately. \square

From Theorem 3.9, we can also extract the following analogue of Lemma 2.1 for double-flips.

Corollary 3.10. *If $\alpha, \alpha'' \in \text{Acyc}(G)$ satisfy $\alpha \approx \alpha''$, then we can reach α'' from α in no more than $4n^4(1 + o(1))$ double-flips.*

Proof. Given an n -vertex graph G and $\alpha, \alpha'' \in \text{Acyc}(G)$ satisfying $\alpha \approx \alpha''$, extract linear extensions $\sigma \in \mathcal{L}(\alpha)$, $\tau \in \mathcal{L}(\alpha'')$, and consider σ and τ as vertices of $\text{FS}(\text{Cycle}_n, G)$. By Theorem 3.9, $d(\sigma, \tau) \leq 8n^4(1 + o(1))$, so let $\Sigma = \{\sigma_i\}_{i=0}^\lambda$ be a shortest sequence of swaps from σ to τ , so $\lambda \leq 8n^4(1 + o(1))$. Let $\Sigma_0 = \{\sigma_{i_j}\}_{j=0}^{\lambda'}$ be the subsequence of Σ consisting of all indices i_j for which $\sigma_{i_{j+1}}$ is reached from σ_{i_j} by a (Cycle_n, G) -friendly swap across the edge $\{1, n\}$. Since λ is smallest possible, two consecutive swaps of Σ cannot both be across the edge $\{1, n\}$, so $\lambda' \leq 4n^4(1 + o(1))$. We will now describe how to use Σ_0 to construct a sequence $\Sigma' = \{\alpha_j\}_{j=0}^{\lambda'+1}$ of acyclic orientations, with $\alpha_0 = \alpha$ and $\alpha_{\lambda'+1} = \alpha''$, for which α_j is reachable from α_{j-1} by a double-flip for all $j \in [\lambda' + 1]$. The desired result will then follow immediately.

Since we reached σ_{i_0} from σ by swapping along the graph $\text{FS}(\text{Path}_n, G)$ (specifically, the copy of Path_n in Cycle_n resulting from excluding the edge $\{1, n\}$), it follows from Theorem 2.4 that $\sigma_{i_0} \in \mathcal{L}(\alpha)$. Let α_1 be the result of taking α and performing a double-flip which involves an inflip on the source $\sigma_{i_0}(1)$ and an outflip on the sink $\sigma_{i_0}(n)$. Note that $\{\sigma_{i_0}(1), \sigma_{i_0}(n)\} \in E(Y)$ (we swapped these two vertices to reach $\sigma_{i_{0+1}}$ from σ_{i_0}), so $\{\sigma_{i_0}(1), \sigma_{i_0}(n)\} \notin E(\bar{Y})$, from which it follows that this is a valid double-flip.³ It is easy to check that $\sigma_{i_{0+1}} \in \mathcal{L}(\alpha_1)$, and by appealing to Theorem 2.4 as before, $\sigma_{i_1} \in \mathcal{L}(\alpha_1)$. Continuing like this sequentially on $j \in [\lambda' + 1]$ (the preceding discussion being the $j = 1$ case) yields the desired sequence Σ' : for the case $j = \lambda' + 1$, it follows as before from Theorem 2.4 that $\sigma_{i_{\lambda'+1}}$ and τ are linear extensions of the poset (i.e., associated to the same acyclic orientation of G), so the final acyclic orientation in Σ' is $\alpha_G(\tau) = \alpha''$. \square

4. Proof of Main Result

We devote this section to answering Question 1.2 in the negative, establishing Theorem 1.3.

4.1. The Graphs X_L and Y_L

We begin with the following observation. One can understand this as the central vertex of Star_n acting as a “knob” rotating around Cycle_n , and all other vertices of $V(\text{Star}_n)$ moving cyclically around it: $n(n-1)$ such swaps in the same direction are needed for all vertices of Star_n to return to their original positions in the starting configuration. This interpretation will help motivate our construction.

Lemma 4.1. *Every connected component of $\text{FS}(\text{Cycle}_n, \text{Star}_n)$ is isomorphic to $\text{Cycle}_{n(n-1)}$.*

³This correspondence between double-flips and paths in $\text{FS}(\text{Cycle}_n, G)$ is the same as that which was observed in the first paragraph of the proof of [DK21, Theorem 4.1].

Proof. Consider a component \mathcal{C} of $\text{FS}(\text{Cycle}_n, \text{Star}_n)$ with permutation $\sigma = \sigma(1) \cdots \sigma(n)$ such that $\sigma(1)$ is the central vertex of Star_n . With $V(\text{Cycle}_{n(n-1)}) = [n(n-1)]$, construct $\varphi : V(\text{Cycle}_{n(n-1)}) \rightarrow V(\mathcal{C})$ by defining $\varphi(i)$ to be the permutation achieved by starting from σ and swapping $\sigma(1)$ rightward i times (e.g., $\varphi(1) = \sigma(2)\sigma(1) \cdots \sigma(n)$). It follows that φ is a graph isomorphism. \square

We will now construct the graphs X_L and Y_L , for every integer $L \geq 1$, that we study to prove Theorem 1.3. In the following description, assume we have fixed some arbitrary integer $L \geq 1$.

The Graph X_L .

The graph X_L contains an $L \times 2$ array of cycle subgraphs, with adjacent cycles intersecting in exactly one vertex. Say X_L has L layers, indexed by $\ell \in [L]$; we will subscript subgraphs and vertices corresponding to the “left column” of X_L by a , and those in the right by b . As such, we denote the left and right cycle subgraphs in layer ℓ by \mathcal{C}_a^ℓ and \mathcal{C}_b^ℓ , respectively. Corresponding to each \mathcal{C}_a^ℓ and \mathcal{C}_b^ℓ is a path subgraph of X_L extending out of it; that corresponding to \mathcal{C}_a^ℓ is denoted \mathcal{P}_a^ℓ , and similarly \mathcal{P}_b^ℓ for \mathcal{C}_b^ℓ . Denote the subgraph of X_L consisting of the ℓ^{th} layer by X^ℓ . The subgraph consisting of \mathcal{P}_a^ℓ and \mathcal{C}_a^ℓ is denoted X_a^ℓ , and similarly X_b^ℓ for \mathcal{P}_b^ℓ and \mathcal{C}_b^ℓ . Denote, whenever they are defined for $\ell \in [L]$,

$$v_a^\ell = V(\mathcal{P}_a^\ell) \cap V(\mathcal{C}_a^\ell), \quad v_b^\ell = V(\mathcal{P}_b^\ell) \cap V(\mathcal{C}_b^\ell), \quad v^\ell = V(\mathcal{C}_a^\ell) \cap V(\mathcal{C}_b^\ell),$$

$$v_a^{\ell, \ell+1} = V(\mathcal{C}_a^\ell) \cap V(\mathcal{C}_a^{\ell+1}), \quad v_b^{\ell, \ell+1} = V(\mathcal{C}_b^\ell) \cap V(\mathcal{C}_b^{\ell+1}).$$

For each of the following sets, we place three inner vertices in the path in \mathcal{C}_a^ℓ between the two vertices in the set:

$$\{v_a^\ell, v_a^{\ell, \ell+1}\}, \{v_a^{\ell, \ell+1}, v^\ell\}, \{v^\ell, v_a^{\ell-1, \ell}\}, \{v_a^{\ell-1, \ell}, v_a^\ell\}.$$

The analogous statement for \mathcal{C}_b^ℓ holds. The exceptions are layers 1 and L : we place seven inner vertices in the upper path from v_a^1 to v^1 in \mathcal{C}_a^1 and the upper path from v_b^1 to v^1 in \mathcal{C}_b^1 , and seven inner vertices in the lower path from v_a^L to v^L in \mathcal{C}_a^L and from v_b^L to v^L in \mathcal{C}_b^L . It follows from our construction that for every $\ell \in [L]$,

$$|V(\mathcal{C}_a^\ell)| = |V(\mathcal{C}_b^\ell)| = 16.$$

We will also set,⁴ for every $\ell \in [L]$,

$$|V(\mathcal{P}_a^\ell)| = 16, \quad |V(\mathcal{P}_b^\ell)| = 15,$$

so that the graph X_L has

$$n = 60 + 58(L - 1) = 58L + 2$$

vertices. (Indeed, it can be checked that layer 1 has 60 vertices, and for each subsequent layer, we add 58 new vertices to the graph Y_L .) Figure 4.1 illustrates this construction for $L = 3$.

⁴It will be important that, for every $\ell \in [L]$, \mathcal{P}_a^ℓ has exactly one more vertex than \mathcal{P}_b^ℓ . The choice of the lengths of these paths, as well as the number of inner vertices in the segments of the cycle subgraphs, is not terribly important as long as they are not too small. The values we chose here suffice.

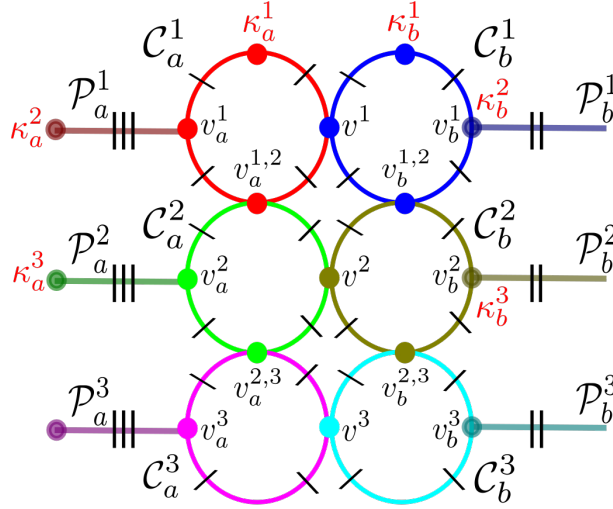


Figure 4.1: Labeled schematic diagram of the construction for X_3 . Subgraphs of X_3 marked a specific color correspond to the σ_s -preimages of the vertices of the same color in Figure 4.2. We take care in appropriately coloring the vertices between two adjacent cycle subgraphs and between adjacent path and cycle subgraphs. Paths marked with one hatch mark have three inner vertices. The paths \mathcal{P}_b^i with two hatch marks have 15 vertices, while paths \mathcal{P}_a^i with three hatch marks have 16 vertices.

The Graph Y_L .

We construct a complementary graph Y_L for each X_L : we assign to each cycle subgraph \mathcal{C}_a^ℓ and \mathcal{C}_b^ℓ of X_L a corresponding “knob vertex” in $V(Y_L)$, denoted κ_a^ℓ and κ_b^ℓ , respectively; we set a collection of vertices of $V(Y_L)$ to swap only with each knob. The construction of Y_L proceeds sequentially according to $\ell \in [L]$. Take two disjoint copies of Star_{15} , denoted \mathcal{S}_a^1 and \mathcal{S}_b^1 , with central vertices κ_a^1 and κ_b^1 , respectively, and a complete bipartite graph \mathcal{K}^1 with 15 vertices in each of its partite sets \mathcal{K}_a^1 and \mathcal{K}_b^1 . Set κ_a^1 and κ_b^1 adjacent to all the vertices in $V(\mathcal{K}^1)$. If $L = 1$, this completes the construction of Y_L . If $L > 1$, take one vertex each in \mathcal{K}_a^1 and \mathcal{K}_b^1 , which shall correspond to κ_a^2 and κ_b^2 , central vertices of star subgraphs (again both isomorphic to Star_{15}) \mathcal{S}_a^2 and \mathcal{S}_b^2 , respectively, and also construct a complete bipartite graph \mathcal{K}^2 with 15 vertices in each of its partite sets \mathcal{K}_a^2 and \mathcal{K}_b^2 . Set κ_a^2 and κ_b^2 adjacent to all the vertices in $V(\mathcal{K}^2)$. Proceed similarly: for $2 \leq \ell \leq L$, take two vertices of $\mathcal{K}^{\ell-1}$ in opposite partite sets and construct \mathcal{S}_a^ℓ , \mathcal{S}_b^ℓ , and \mathcal{K}^ℓ , related as before, until all $n = 58L + 2$ vertices are exhausted. We shall often refer to vertices κ_a^ℓ and κ_b^ℓ as *knob vertices* of Y_L . Figure 4.2 illustrates this construction for $L = 3$, while Figure 4.3 provides a “collapsed” view of our construction.

The Starting Configuration σ_s and its Connected Component \mathcal{C} .

Take an arbitrary $L \geq 1$ and graphs X_L, Y_L . We are now going to describe a specific starting configuration $\sigma_s(X_L, Y_L) \in V(\text{FS}(X_L, Y_L))$ which lies in the connected component $\mathcal{C}(X_L, Y_L)$ of $\text{FS}(X_L, Y_L)$; we will later show that there exists a different configuration in $\mathcal{C}(X_L, Y_L)$ whose




Figure 4.2: Labeled schematic diagram of the construction for Y_3 . The vertices of Y_3 marked with a particular color correspond to the σ_s -images of the vertices of the same color in Figure 4.1.

distance from $\sigma_s(X_L, Y_L)$ is $e^{\Omega(n)}$. Henceforth, we abbreviate $\sigma_s(X_L, Y_L)$ and $\mathcal{C}(X_L, Y_L)$ to σ_s and \mathcal{C} . In forthcoming discussions, X_L and Y_L will be understood to be arbitrary such graphs on the same number of vertices.

Take all 15 vertices in $V(\mathcal{K}_a^1)$ and place them onto $V(\mathcal{P}_a^1) \setminus \{v_a^1\}$, and the 15 vertices in $V(\mathcal{K}_b^1)$ onto $V(\mathcal{P}_b^1)$; if $L > 1$, we place κ_a^2 onto the leftmost vertex of $V(\mathcal{P}_a^1)$ and κ_b^2 onto v_b^1 . Now take subgraph \mathcal{S}_a^1 of Y_L : place κ_a^1 onto the middle vertex of the upper path between v_a^1 and v^1 (which has seven vertices), and place all 14 leaves of \mathcal{S}_a^1 onto the remaining 14 vertices of $V(\mathcal{C}_a^1) \setminus \{v^1\}$ in some way. Similarly, take \mathcal{S}_b^1 : place κ_b^1 onto the middle vertex of the upper path between v^1 and v_b^1 , and place all 14 leaves of \mathcal{S}_b^1 onto the remaining 14 vertices of $V(\mathcal{C}_b^1)$. This has filled all mappings on the subgraph $V(X^1)$ of X_L by vertices in $V(\mathcal{K}^1)$, $V(\mathcal{S}_a^1)$, and $V(\mathcal{S}_b^1)$, and thus yields σ_s if $L = 1$.

Proceed sequentially according to the layer $\ell \in [L]$: say we placed all vertices of $V(\mathcal{K}^i)$, $V(\mathcal{S}_a^i)$, and $V(\mathcal{S}_b^i)$ for $i < \ell$ onto the corresponding $V(X^i)$ of X_L . Place all 15 vertices in $V(\mathcal{K}_a^\ell)$ onto $V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}$, and the 15 vertices in $V(\mathcal{K}_b^\ell)$ onto $V(\mathcal{P}_b^\ell)$; if $L > \ell$, place $\kappa_a^{\ell+1}$ onto the leftmost vertex of $V(\mathcal{P}_a^\ell)$ and $\kappa_b^{\ell+1}$ onto v_b^ℓ . Now take \mathcal{S}_a^ℓ , and place its 14 leaves onto the remaining 14 vertices in $V(\mathcal{C}_a^\ell) \setminus \{v^\ell\}$. Similarly take \mathcal{S}_b^ℓ , and place its 14 leaves onto the 14 remaining vertices in $V(\mathcal{C}_b^\ell)$. An illustration of this starting configuration is given in Figures 4.1 and 4.2: the vertices of a particular color in Figure 4.2 are placed upon the correspondingly colored subgraph in Figure 4.1 to achieve $\sigma_s \in V(\text{FS}(X_L, Y_L))$.

Remark 4.2. By the construction of $\sigma_s \in V(\text{FS}(X_L, Y_L))$, for any $\ell \in [L]$,

$$V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\} \subset \sigma_s(V(\mathcal{C}_a^\ell)), \quad V(\mathcal{S}_b^\ell) \setminus \{\kappa_b^\ell\} \subset \sigma_s(V(\mathcal{C}_b^\ell)).$$




Figure 4.3: A simplified schematic diagram of Y_4 to illustrate the neighborhoods of different kinds of vertices. Here, all subgraphs \mathcal{S}_a^ℓ , \mathcal{S}_b^ℓ , \mathcal{K}_a^ℓ , and \mathcal{K}_b^ℓ are to be understood as excluding any knob vertices.

As such, all leaves of a star subgraph \mathcal{S}_a^ℓ or \mathcal{S}_b^ℓ of Y_L are placed onto a corresponding cycle subgraph \mathcal{C}_a^ℓ or \mathcal{C}_b^ℓ of X_L , respectively. This yields that, for any $\ell \in [L]$,

$$|\sigma_s(V(\mathcal{C}_a^\ell) \setminus (V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\}))| = |\sigma_s(V(\mathcal{C}_b^\ell) \setminus (V(\mathcal{S}_b^\ell) \setminus \{\kappa_b^\ell\}))| = 2.$$

In other words, the number of vertices upon any cycle subgraph \mathcal{C}_a^ℓ or \mathcal{C}_b^ℓ of X_L which are not leaves of the corresponding star subgraph of Y_L , under σ_s , is exactly two.

We introduce the following definition for notational convenience in forthcoming arguments.

Definition 4.3. Fix $\ell \in [L]$.

- The *boundary* $\text{bd}(\mathcal{C}_a^\ell)$ of \mathcal{C}_a^ℓ is the subset of $\{v_a^\ell, v_a^{\ell-1, \ell}, v_a^{\ell, \ell+1}, v^\ell\}$ defined for ℓ .
- The *boundary* $\text{bd}(\mathcal{C}_b^\ell)$ of \mathcal{C}_b^ℓ is the subset of $\{v_b^\ell, v_b^{\ell-1, \ell}, v_b^{\ell, \ell+1}, v^\ell\}$ defined for ℓ .

In Subsections 4.2 and 4.3, unless otherwise stated, we fix an arbitrary integer $L \geq 1$ and refer to the graphs X_L and Y_L , with σ_s denoting the corresponding starting configuration. We elect to refer to paths in $\text{FS}(X_L, Y_L)$ as *swap sequences*, which are denoted by the vertices and edges in $\text{FS}(X_L, Y_L)$ that constitute the path. More specifically, a swap sequence of length λ is a sequence of vertices $\Sigma = \{\sigma_i\}_{i=0}^\lambda \subseteq V(\text{FS}(X_L, Y_L))$ for which $\{\sigma_{i-1}, \sigma_i\} \in E(\text{FS}(X_L, Y_L))$ for all $i \in [\lambda]$.

4.2. Configurations in \mathcal{C}

In this subsection, we derive properties satisfied by all vertices in \mathcal{C} . Intuitively, our aim in this subsection is to uncover many conditions satisfied by all of the vertices in \mathcal{C} , which has the effect of producing strong rigidities on the corresponding swapping problem. These rigidities will allow us to argue in Subsection 4.3 that in order to move certain vertices in Y_L down and across the graph X_L , we necessarily must perform very specific sequences of swaps.

Remark 4.2 observes that in the starting configuration σ_s , the leaves of any star graph \mathcal{S}_a^ℓ or \mathcal{S}_b^ℓ lie upon the vertices of \mathcal{C}_a^ℓ and \mathcal{C}_b^ℓ , respectively. In particular, for any cycle subgraph \mathcal{C}_a^ℓ in X_L , exactly two vertices that are not leaves of \mathcal{S}_a^ℓ lie upon them; an analogous statement holds for cycle subgraphs of the form \mathcal{C}_b^ℓ . We begin our study of \mathcal{C} by establishing that this property is maintained after any sequence of swaps in $\text{FS}(X_L, Y_L)$ beginning at σ_s , i.e., that all vertices in \mathcal{C} satisfy this property: we prove this in Proposition 4.4.

Proposition 4.4. *Any $\sigma \in V(\mathcal{C})$ satisfies, for all $\ell \in [L]$,*

$$V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\} \subset \sigma(V(\mathcal{C}_a^\ell)) \text{ and } V(\mathcal{S}_b^\ell) \setminus \{\kappa_b^\ell\} \subset \sigma(V(\mathcal{C}_b^\ell)).$$

As in Remark 4.2, this means that for any cycle subgraph \mathcal{C}_a^ℓ or \mathcal{C}_b^ℓ in X_L and $\sigma \in V(\mathcal{C})$,

$$|\sigma(V(\mathcal{C}_a^\ell)) \setminus (V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\})| = |\sigma(V(\mathcal{C}_b^\ell)) \setminus (V(\mathcal{S}_b^\ell) \setminus \{\kappa_b^\ell\})| = 2,$$

since $|V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\}| = |V(\mathcal{S}_b^\ell) \setminus \{\kappa_b^\ell\}| = 14$, and $|V(\mathcal{C}_a^\ell)| = |V(\mathcal{C}_b^\ell)| = 16$ for all $\ell \in [L]$.

Remark 4.5. Although Proposition 4.4 describes a global property maintained by all configurations in \mathcal{C} , we frequently appeal to it (for sake of brevity) as a local property satisfied by specific configurations in \mathcal{C} during the proof of Proposition 4.4.⁵ This practice of localizing a more global statement to a particular configuration will also be utilized for other results in later proofs in this section, and it should not raise any ambiguity whenever it is invoked.

Proof of Proposition 4.4. Assume (towards a contradiction) that the proposition is false, so there exists a swap sequence $\Sigma = \{\sigma_i\}_{i=0}^\lambda$ with $\sigma_0 = \sigma_s$ in \mathcal{C} of shortest length λ containing a vertex violating Proposition 4.4: σ_λ violates Proposition 4.4, while all σ_i for $i < \lambda$ satisfy it, and $\lambda \geq 1$. Thus, there exists a star subgraph \mathcal{S} (of form \mathcal{S}_a^ℓ or \mathcal{S}_b^ℓ) of Y_L and a leaf $\mu \in V(\mathcal{S})$ such that $\sigma_{\lambda-1}^{-1}(\mu)$ is in the appropriate cycle subgraph, but $\sigma_\lambda^{-1}(\mu)$ is not. Say $\mathcal{S} = \mathcal{S}_a^\ell$ for $\ell \in [L]$: raising a contradiction when $\mathcal{S} = \mathcal{S}_b^\ell$ is entirely analogous. Here, $\mu \in V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\}$ has $N_{Y_L}(\mu) = \{\kappa_a^\ell\}$ and $\sigma_{\lambda-1}^{-1}(\mu) \in V(\mathcal{C}_a^\ell)$, $\sigma_\lambda^{-1}(\mu) \notin V(\mathcal{C}_a^\ell)$, so $\sigma_{\lambda-1}^{-1}(\mu) \in \text{bd}(\mathcal{C}_a^\ell)$ and σ_λ is reached from $\sigma_{\lambda-1}$ by swapping μ and κ_a^ℓ . Figure 4.4 depicts the configurations described in the following two cases.

Case 1: We have $\sigma_{\lambda-1}^{-1}(\mu) = v_a^\ell$. Here, $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell) \in N_{X_L}(v_a^\ell) \cap V(\mathcal{P}_a^\ell)$. Let $\xi < \lambda - 1$ be the final such index with $\sigma_\xi^{-1}(\kappa_a^\ell) \notin V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}$; ξ is well-defined since

$$\sigma_s^{-1}(\kappa_a^\ell) \notin V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\},$$

⁵Making this clarification is important, as the proof proceeds by assuming (towards a contradiction) that Proposition 4.4 is satisfied by particular configurations in \mathcal{C} and is violated by another.

which implies $\sigma_s \neq \sigma_{\lambda-1}$, so $\lambda \geq 2$. By the definition of ξ and $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell) \in N_{X_L}(v_a^\ell) \cap V(\mathcal{P}_a^\ell)$,

$$\sigma_j^{-1}(\kappa_a^\ell) \in V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\} \text{ for } \xi + 1 \leq j \leq \lambda - 1. \quad (4.1)$$

Necessarily, $\sigma_\xi^{-1}(\kappa_a^\ell) = v_a^\ell$ and $\sigma_{\xi+1}^{-1}(\kappa_a^\ell) \in N_{X_L}(v_a^\ell) \cap V(\mathcal{P}_a^\ell)$, so

$$\sigma_\xi^{-1}(\mu) = \sigma_{\xi+1}^{-1}(\mu) \in V(\mathcal{C}_a^\ell) \setminus \{v_a^\ell\};$$

note that σ_ξ satisfies Proposition 4.4. Since $N_{Y_L}(\mu) = \{\kappa_a^\ell\}$ and there are no edges between $V(\mathcal{C}_a^\ell) \setminus \{v_a^\ell\}$ and $V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}$, it follows from (4.1) that $\sigma_j^{-1}(\mu)$ is fixed for $\xi \leq j \leq \lambda - 1$, so

$$\sigma_\xi^{-1}(\mu) = \sigma_{\lambda-1}^{-1}(\mu) \in V(\mathcal{C}_a^\ell) \setminus \{v_a^\ell\},$$

contradicting $\sigma_{\lambda-1}^{-1}(\mu) = v_a^\ell$.

Case 2: We have $\sigma_{\lambda-1}^{-1}(\mu) \neq v_a^\ell$. Here, $\sigma_{\lambda-1}^{-1}(\mu) \in \text{bd}(\mathcal{C}_a^\ell) \setminus \{v_a^\ell\}$, and

$$\sigma_{\lambda-1}^{-1}(\kappa_a^\ell) = \sigma_{\lambda-1}^{-1}(\mu) \in N_{X_L}(\sigma_{\lambda-1}^{-1}(\mu)) \setminus V(\mathcal{C}_a^\ell).$$

Proceeding backwards in Σ , it must be that either

$$\sigma_{\lambda-2}^{-1}(\mu) \neq \sigma_{\lambda-1}^{-1}(\mu) \text{ or } \sigma_{\lambda-2}^{-1}(\kappa_a^\ell) \neq \sigma_{\lambda-1}^{-1}(\kappa_a^\ell);$$

note that $\lambda \geq 2$, since $\sigma_s^{-1}(\kappa_a^\ell) \neq \sigma_{\lambda-1}^{-1}(\kappa_a^\ell)$. Indeed, if not, swapping μ and κ_a^ℓ directly from $\sigma_{\lambda-2}$ raises a contradiction on λ being minimal. Now, $N_{Y_L}(\mu) = \{\kappa_a^\ell\}$ implies

$$\sigma_{\lambda-2}^{-1}(\kappa_a^\ell) \neq \sigma_{\lambda-1}^{-1}(\kappa_a^\ell) \text{ and } \sigma_{\lambda-2}^{-1}(\mu) = \sigma_{\lambda-1}^{-1}(\mu),$$

since if both preimages differ, $\sigma_{\lambda-2} = \sigma_\lambda$. Thus, $\sigma_{\lambda-2}^{-1}(\kappa_a^\ell) \notin V(\mathcal{C}_a^\ell)$ and

$$\sigma_{\lambda-2}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) \notin V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\}$$

by Proposition 4.4 (on $\sigma_{\lambda-2}$), so $\sigma_{\lambda-2}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell))$ is not a leaf (see $N_{Y_L}(\kappa_a^\ell)$). We further assume that

$$\sigma_{\lambda-2}^{-1}(\mu) = \sigma_{\lambda-1}^{-1}(\mu) = v_a^\ell;$$

raising a contradiction when this vertex is $v_a^{\ell-1, \ell}$ or $v_a^{\ell, \ell+1}$ can be done analogously. Therefore

$$\sigma_{\lambda-2}^{-1}(\{\kappa_a^\ell, \mu\}) \subset V(\mathcal{C}_b^\ell) \text{ and } \sigma_{\lambda-2}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) \in V(\mathcal{C}_b^\ell).$$

Altogether, we have that

$$|\sigma_{\lambda-2}(V(\mathcal{C}_b^\ell)) \setminus (V(\mathcal{S}_b^\ell) \setminus \{\kappa_b^\ell\})| \geq 3,$$

and since $|V(\mathcal{C}_b^\ell)| = 16$ and $|V(\mathcal{S}_b^\ell) \setminus \{\kappa_b^\ell\}| = 14$, $V(\mathcal{S}_b^\ell) \setminus \{\kappa_b^\ell\} \not\subset \sigma_{\lambda-2}(V(\mathcal{C}_b^\ell))$. Thus, $\sigma_{\lambda-2}$ violates Proposition 4.4, contradicting λ being minimal. \square




Figure 4.4: Configurations in Σ raising a contradiction for both cases in the proof of Proposition 4.4.

Proposition 4.4 restricts the preimages of the leaves of \mathcal{S}_a^ℓ and \mathcal{S}_b^ℓ under any $\sigma \in V(\mathcal{C})$. We now derive a restriction on the preimages of all other vertices in $V(Y_L)$ under any $\sigma \in V(\mathcal{C})$. As Proposition 4.6 formalizes, for such σ , any vertex in $V(Y_L)$ is close to its preimage in σ_s .

Proposition 4.6. *Any configuration $\sigma \in V(\mathcal{C})$ must satisfy the following four properties.*

1. *The layer 1 knob vertices lie upon the corresponding subgraph of X^1 , i.e.,*

$$\sigma^{-1}(\kappa_a^1) \in V(X_a^1) \text{ and } \sigma^{-1}(\kappa_b^1) \in V(X_b^1).$$

2. *For $2 \leq \ell \leq L$, the layer ℓ knob vertices lie upon the subgraphs $X^{\ell-1}$ or X^ℓ , i.e.,*

$$\{\sigma^{-1}(\kappa_a^\ell), \sigma^{-1}(\kappa_b^\ell)\} \subset V(X^{\ell-1}) \cup V(X^\ell).$$

3. *For $\ell \in [L-1]$, any vertex in $V(\mathcal{K}^\ell)$ that is not a layer $\ell+1$ knob lies upon X^ℓ , i.e.,*

$$\sigma^{-1}(V(\mathcal{K}^\ell) \setminus \{\kappa_a^{\ell+1}, \kappa_b^{\ell+1}\}) \subset V(X^\ell),$$

and every vertex in $V(\mathcal{K}^L)$ lies upon X^L , i.e.,

$$\sigma^{-1}(V(\mathcal{K}^L)) \subset V(X^L).$$

4. *For $\ell \in [L]$, there is at most one $\mu \in V(\mathcal{K}^\ell)$ not in $V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$, i.e.,*

$$|\sigma^{-1}(V(\mathcal{K}^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| \leq 1.$$

Confirming that the starting configuration σ_s satisfies these four properties is straightforward. Case 4 of the proof of Proposition 4.6 relies on the following Lemma 4.7, which is illustrated in Figure 4.5. In the statement of the lemma, we elect to index the final term of the swap sequence by $\lambda-1$ as this is where the result applies in the proof of Proposition 4.6.


Lemma 4.7. *Let $\Sigma = \{\sigma_i\}_{i=0}^{\lambda-1}$ with $\sigma_0 = \sigma_s$, $\lambda \geq 1$ be a swap sequence in \mathcal{C} such that for all $1 \leq i \leq \lambda-1$, σ_i satisfies the four properties of Proposition 4.6. Then for all $0 \leq i \leq \lambda-1$ and $\ell \in [L]$, the following two statements hold.*

1. If $\sigma_i^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$, then

$$\sigma_i(\{v_a^\ell, v_b^\ell\}) \subset V(\mathcal{K}^\ell) \implies |\sigma_i^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap ((V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}) \cup (V(\mathcal{P}_b^\ell) \setminus \{v_b^\ell\}))| = 1.$$

2. If $\sigma_i^{-1}(V(\mathcal{K}^\ell)) \not\subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$, then

$$\begin{aligned} \sigma_i(v_a^\ell) \in V(\mathcal{K}^\ell) &\implies \sigma_i^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap (V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}) \neq \emptyset, \\ \sigma_i(v_b^\ell) \in V(\mathcal{K}^\ell) &\implies \sigma_i^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap (V(\mathcal{P}_b^\ell) \setminus \{v_b^\ell\}) \neq \emptyset. \end{aligned}$$



(a) Lemma 4.7(1) on the configuration σ_i . If $\sigma_i^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$ and $\sigma_i(\{v_a^\ell, v_b^\ell\}) \subset V(\mathcal{K}^\ell)$, then either κ_a^ℓ or κ_b^ℓ lies upon $V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}$ or $V(\mathcal{P}_b^\ell) \setminus \{v_b^\ell\}$.

(b) First implication of Lemma 4.7(2) on the configuration σ_i . If there exists some $\mu \in V(\mathcal{K}^\ell)$ with $\sigma_i^{-1}(\mu) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$, and $\sigma_i(v_a^\ell) \in V(\mathcal{K}^\ell)$, then either κ_a^ℓ or κ_b^ℓ lies upon $V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}$.

Figure 4.5: Illustrations for both parts of Lemma 4.7 for some $\sigma_i \in \Sigma$. Subgraphs/vertices colored in red correspond to σ_i -preimages of $V(\mathcal{K}^\ell)$, while σ_i -preimages of elements in $\{\kappa_a^\ell, \kappa_b^\ell\}$ are colored in blue. For Figure 4.5b, note that by appealing to Proposition 4.6(4) and comparing cardinalities, we can deduce that at most two vertices of $\sigma_i(V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))$ can fail to lie in $V(\mathcal{K}^\ell)$.

Proof of Lemma 4.7. Fix $\Sigma = \{\sigma_i\}_{i=0}^{\lambda-1}$ to be a swap sequence satisfying the assumptions of Lemma 4.7. We prove the two statements of Lemma 4.7 hold for all $\ell \in [L]$ inductively for $0 \leq i \leq \lambda - 1$. They can be checked to hold for all $\ell \in [L]$ when $i = 0$, so assume they are true for some $0 \leq i < \lambda - 1$. We prove that σ_{i+1} satisfies both statements for all $\ell \in [L]$. In what follows, assume we refer (unless stated otherwise) to some fixed, arbitrary $\ell \in [L]$. We break into cases based on whether or not $\sigma_i^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$.

Case 1: We have $\sigma_i^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$. We will further break into subcases based on whether or not $\sigma_i(\{v_a^\ell, v_b^\ell\}) \subset V(\mathcal{K}^\ell)$.

Subcase 1.1: We have $\sigma_i(\{v_a^\ell, v_b^\ell\}) \subset V(\mathcal{K}^\ell)$. By the induction hypothesis, we have that

$$|\sigma_i^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap ((V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}) \cup (V(\mathcal{P}_b^\ell) \setminus \{v_b^\ell\}))| = 1. \quad (4.2)$$

If $\sigma_i(v_a^\ell) = \sigma_{i+1}(v_a^\ell)$ and $\sigma_i(v_b^\ell) = \sigma_{i+1}(v_b^\ell)$, then σ_{i+1} satisfies Lemma 4.7(1) since

$$\sigma_{i+1}^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell), \quad |\sigma_{i+1}^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap ((V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}) \cup (V(\mathcal{P}_b^\ell) \setminus \{v_b^\ell\}))| = 1,$$

and satisfies Lemma 4.7(2) trivially.⁶ So consider the setting where either $\sigma_i(v_a^\ell) \neq \sigma_{i+1}(v_a^\ell)$ or $\sigma_i(v_b^\ell) \neq \sigma_{i+1}(v_b^\ell)$: say $\sigma_i(v_a^\ell) \neq \sigma_{i+1}(v_a^\ell)$ (the setting $\sigma_i(v_b^\ell) \neq \sigma_{i+1}(v_b^\ell)$ is analogous). If

$$\sigma_{i+1}^{-1}(\sigma_i(v_a^\ell)) \in N_{X_L}(v_a^\ell) \cap V(\mathcal{P}_a^\ell),$$

then σ_{i+1} satisfies Lemma 4.7(1). Indeed, since $\kappa_a^\ell, \kappa_b^\ell \notin V(\mathcal{K}^\ell)$, the hypothesis $\sigma_{i+1}(\{v_a^\ell, v_b^\ell\}) \subset V(\mathcal{K}^\ell)$ implies that $\sigma_i^{-1}(\kappa_a^\ell) = \sigma_{i+1}^{-1}(\kappa_a^\ell)$ and $\sigma_i^{-1}(\kappa_b^\ell) = \sigma_{i+1}^{-1}(\kappa_b^\ell)$. If

$$\sigma_{i+1}^{-1}(\sigma_i(v_a^\ell)) \notin V(\mathcal{P}_a^\ell),$$

then since $\sigma_i(v_a^\ell) \in V(\mathcal{K}^\ell)$, we have that

$$\sigma_{i+1}^{-1}(V(\mathcal{K}^\ell)) \not\subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell).$$

From studying the neighborhoods of vertices in $V(\mathcal{K}^\ell)$ to produce possibilities for $\sigma_{i+1}(v_a^\ell)$, Propositions 4.4 and 4.6(2,3)⁷ imply

$$\sigma_{i+1}(v_a^\ell) \in V(\mathcal{K}^\ell) \cup \{\kappa_a^\ell, \kappa_b^\ell\}$$

(consider the possible vertices in $N_{Y_L}(\sigma_i(v_a^\ell))$), from which $\sigma_i^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$ implies

$$\sigma_{i+1}(v_a^\ell) \in \{\kappa_a^\ell, \kappa_b^\ell\}.$$

This yields $\ell = 1$. Indeed, if $\ell \geq 2$, then σ_i violates Proposition 4.6(4) on layer $\ell - 1$, since with (4.2),

$$\sigma_i^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \subset \sigma_i^{-1}(V(\mathcal{K}^{\ell-1})) \setminus (V(\mathcal{P}_a^{\ell-1}) \cup V(\mathcal{P}_b^{\ell-1})),$$

so that $\sigma_{i+1}(v_a^1) = \kappa_a^1$ by Proposition 4.6(1) on σ_{i+1} . This result, with Proposition 4.6(1) (on σ_i) and (4.2), yields

$$\sigma_i^{-1}(\kappa_b^1) = \sigma_{i+1}^{-1}(\kappa_b^1) \in V(\mathcal{P}_b^1) \setminus \{v_b^1\},$$

so σ_{i+1} satisfies Lemma 4.7(2).

Subcase 1.2: We have $\sigma_i(\{v_a^\ell, v_b^\ell\}) \not\subset V(\mathcal{K}^\ell)$. Since $|V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)| = 31$ and $|V(\mathcal{K}^\ell)| = 30$, and recalling our initial assumption $\sigma_i^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$, we have that

$$|\sigma_i(\{v_a^\ell, v_b^\ell\}) \cap V(\mathcal{K}^\ell)| = 1, \quad \sigma_i((V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}) \cup (V(\mathcal{P}_b^\ell) \setminus \{v_b^\ell\})) \subset V(\mathcal{K}^\ell). \quad (4.3)$$

⁶Generally, in what follows, we do not comment on the ‘‘other statement’’ in Lemma 4.7 holding trivially, and only check the statement which applies, depending on whether $\sigma_{i+1}^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$ or not in the given context.

⁷Indeed, σ_{i+1} would violate Proposition 4.6(2) if $\sigma_{i+1}(v_a^\ell) \in \{\kappa_a^{\ell+2}, \kappa_b^{\ell+2}\}$ and Proposition 4.6(3) if $\sigma_{i+1}(v_a^\ell) \in V(\mathcal{K}^{\ell+1}) \setminus \{\kappa_a^{\ell+2}, \kappa_b^{\ell+2}\}$. Henceforth, we do not explicitly make such further distinctions when appealing to multiple properties from Proposition 4.6 together.

Say $\sigma_i(v_a^\ell) \in V(\mathcal{K}^\ell)$; the setting $\sigma_i(v_b^\ell) \in V(\mathcal{K}^\ell)$ is argued analogously. By (4.3), $\sigma_i(v_b^\ell) \notin V(\mathcal{K}^\ell)$. If $\sigma_{i+1}^{-1}(\sigma_i(v_a^\ell)) \notin V(\mathcal{P}_a^\ell)$, then $\sigma_i^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$ and $\sigma_i(v_b^\ell) = \sigma_{i+1}(v_b^\ell) \notin V(\mathcal{K}^\ell)$, implying

$$\sigma_{i+1}^{-1}(V(\mathcal{K}^\ell)) \not\subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell) \text{ and } \sigma_{i+1}(v_a^\ell) \notin V(\mathcal{K}^\ell),$$

so σ_{i+1} satisfies Lemma 4.7(2). If $\sigma_{i+1}^{-1}(\sigma_i(v_a^\ell)) \in V(\mathcal{P}_a^\ell)$, then $\sigma_i^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$ and $\sigma_i(v_b^\ell) \notin V(\mathcal{K}^\ell)$ yield

$$\sigma_{i+1}^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell). \quad (4.4)$$

Studying the neighborhoods of vertices in $V(\mathcal{K}^\ell)$ and recalling that $\sigma_i(v_b^\ell) \notin V(\mathcal{K}^\ell)$ yields that the only way we can have that $\sigma_{i+1}(v_b^\ell) \in V(\mathcal{K}^\ell)$ (exactly when σ_{i+1} does not trivially satisfy Lemma 4.7(1)) without σ_{i+1} violating Proposition 4.4 or 4.6(2,3) is if

$$\sigma_i(v_b^\ell) \in \{\kappa_a^\ell, \kappa_b^\ell\} \text{ and } \sigma_{i+1}^{-1}(\sigma_i(v_b^\ell)) \in N_{X_L}(v_b^\ell) \cap V(\mathcal{P}_b^\ell).$$

These results, along with (4.4), $|V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)| = 31$, $|V(\mathcal{K}^\ell)| = 30$, and $\kappa_a^\ell, \kappa_b^\ell \notin V(\mathcal{K}^\ell)$, imply

$$|\sigma_{i+1}^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap ((V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}) \cup (V(\mathcal{P}_b^\ell) \setminus \{v_b^\ell\}))| = 1,$$

so σ_{i+1} satisfies Lemma 4.7(1).

Case 2: We have $\sigma_i^{-1}(V(\mathcal{K}^\ell)) \not\subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$. By Proposition 4.6(4) (on σ_i), there exists a unique $\mu \in V(\mathcal{K}^\ell)$ such that $\sigma_i^{-1}(\mu) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$, so $|V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)| = 31$ and $|V(\mathcal{K}^\ell)| = 30$ yield

$$|(V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)) \setminus \sigma_i^{-1}(V(\mathcal{K}^\ell))| = 2. \quad (4.5)$$

We break into subcases based on the subset of $\{\sigma_i(v_a^\ell), \sigma_i(v_b^\ell)\}$ that is in $V(\mathcal{K}^\ell)$.

Subcase 2.1: We have $\sigma_i(\{v_a^\ell, v_b^\ell\}) \subset V(\mathcal{K}^\ell)$. By the induction hypothesis,⁸

$$\sigma_i^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap (V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}) \neq \emptyset, \quad \sigma_i^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap (V(\mathcal{P}_b^\ell) \setminus \{v_b^\ell\}) \neq \emptyset.$$

If $\sigma_i(v_a^\ell) = \sigma_{i+1}(v_a^\ell)$ and $\sigma_i(v_b^\ell) = \sigma_{i+1}(v_b^\ell)$, then σ_{i+1} satisfies Lemma 4.7(2). If $\sigma_i(v_a^\ell) \neq \sigma_{i+1}(v_a^\ell)$ (the setting $\sigma_i(v_b^\ell) \neq \sigma_{i+1}(v_b^\ell)$ is argued analogously), we must have that (exactly) one of

$$\sigma_{i+1}^{-1}(\sigma_i(v_a^\ell)) \in N_{X_L}(v_a^\ell) \cap V(\mathcal{P}_a^\ell), \quad \sigma_{i+1}^{-1}(\sigma_i(v_a^\ell)) \in N_{X_L}(v_a^\ell) \setminus V(\mathcal{P}_a^\ell) \text{ and } \sigma_{i+1}(v_a^\ell) = \mu$$

must hold, since σ_{i+1} would otherwise violate Proposition 4.6(4), due to

$$\{\mu, \sigma_i(v_a^\ell)\} \subset \sigma_{i+1}^{-1}(V(\mathcal{K}^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)).$$

As before, σ_{i+1} satisfies Lemma 4.7(2) in both situations.

⁸The resulting observations are enough to deduce that $\ell = 1$, but this is not necessary for the proceeding argument.

Subcase 2.2: We have $\{\sigma_i(v_a^\ell), \sigma_i(v_b^\ell)\} \cap V(\mathcal{K}^\ell) = \{\sigma_i(v_a^\ell)\}$. The setting $\{\sigma_i(v_a^\ell), \sigma_i(v_b^\ell)\} \cap V(\mathcal{K}^\ell) = \{\sigma_i(v_b^\ell)\}$ is argued analogously. The induction hypothesis yields

$$\sigma_i^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap (V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}) \neq \emptyset.$$

From (4.5), we deduce that

$$|\sigma_i^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap (V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\})| = 1 \tag{4.6}$$

and also that

$$(V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)) \setminus \sigma_i^{-1}(V(\mathcal{K}^\ell)) = \{v_b^\ell\} \cup (\sigma_i^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap (V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\})).$$

We can argue as in Subcase 2.1 to deduce that σ_{i+1} satisfies Lemma 4.7(2) if $\sigma_i(v_a^\ell) = \sigma_{i+1}(v_a^\ell)$ and $\sigma_i(v_b^\ell) = \sigma_{i+1}(v_b^\ell)$, or if $\sigma_i(v_a^\ell) \neq \sigma_{i+1}(v_a^\ell)$. If $\sigma_i(v_b^\ell) \neq \sigma_{i+1}(v_b^\ell)$,

$$\sigma_i^{-1}(\mu) = \sigma_{i+1}^{-1}(\mu) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell),$$

so σ_{i+1} satisfies Lemma 4.7(2) if $\sigma_{i+1}(v_b^\ell) \notin V(\mathcal{K}^\ell)$. Thus, assume $\sigma_{i+1}(v_b^\ell) \in V(\mathcal{K}^\ell)$. Studying the neighborhoods of vertices in $V(\mathcal{K}^\ell)$ yields that $\sigma_i(v_b^\ell) \in \{\kappa_a^\ell, \kappa_b^\ell\}$, as σ_i would violate Proposition 4.4 or 4.6(2,3) otherwise. If

$$\sigma_{i+1}^{-1}(\sigma_i(v_b^\ell)) \in N_{X_L}(v_b^\ell) \cap V(\mathcal{P}_b^\ell),$$

σ_{i+1} satisfies Lemma 4.7(2). If

$$\sigma_{i+1}^{-1}(\sigma_i(v_b^\ell)) \in N_{X_L}(v_b^\ell) \setminus V(\mathcal{P}_b^\ell),$$

it must be that $\sigma_{i+1}(v_b^\ell) = \mu$ (recall that $\mu \in V(\mathcal{K}^\ell)$ is the unique such vertex for which $\sigma_i^{-1}(\mu) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$). By (4.6), alongside $\sigma_i(v_b^\ell) \in \{\kappa_a^\ell, \kappa_b^\ell\}$ and $\sigma_{i+1}^{-1}(\sigma_i(v_b^\ell)) \in N_{X_L}(v_b^\ell) \cap V(\mathcal{P}_b^\ell)$, σ_{i+1} satisfies Lemma 4.7(1).

Subcase 2.3: We have $\sigma_i(\{v_a^\ell, v_b^\ell\}) \cap V(\mathcal{K}^\ell) = \emptyset$. From (4.5), we have that

$$\sigma_i^{-1}(V(\mathcal{K}^\ell) \setminus \{\mu\}) = (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)) \setminus \{v_a^\ell, v_b^\ell\}$$

since the LHS is a subset of the RHS and their cardinalities are equal. If $\sigma_{i+1}^{-1}(\mu) \in \{v_a^\ell, v_b^\ell\}$, then

$$\sigma_{i+1}^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell) \text{ and } \sigma_{i+1}(\{v_a^\ell, v_b^\ell\}) \not\subset V(\mathcal{K}^\ell),$$

so σ_{i+1} satisfies Lemma 4.7(1). Now assume $\sigma_{i+1}^{-1}(\mu) \notin \{v_a^\ell, v_b^\ell\}$, from which it easily follows that

$$\sigma_{i+1}^{-1}(V(\mathcal{K}^\ell)) \not\subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell) \text{ and } |\{\sigma_{i+1}(v_a^\ell), \sigma_{i+1}(v_b^\ell)\} \cap V(\mathcal{K}^\ell)| \leq 1.$$

Of course, σ_{i+1} satisfies Lemma 4.7(2) if

$$\{\sigma_{i+1}(v_a^\ell), \sigma_{i+1}(v_b^\ell)\} \cap V(\mathcal{K}^\ell) = \emptyset.$$

If $\sigma_{i+1}(v_a^\ell) \in V(\mathcal{K}^\ell)$ (the setting $\sigma_{i+1}(v_b^\ell) \in V(\mathcal{K}^\ell)$ is argued analogously), then by studying the neighborhoods of vertices in $V(\mathcal{K}^\ell)$, it must be that $\sigma_i(v_a^\ell) \in \{\kappa_a^\ell, \kappa_b^\ell\}$, since σ_i would otherwise violate Proposition 4.4 or 4.6(2,3). Furthermore,

$$\sigma_{i+1}^{-1}(\sigma_i(v_a^\ell)) \in N_{X_L}(v_a^\ell) \cap V(\mathcal{P}_a^\ell),$$

since if $\sigma_{i+1}^{-1}(\sigma_i(v_a^\ell)) \in N_{X_L}(v_a^\ell) \setminus V(\mathcal{P}_a^\ell)$, we would have

$$\{\sigma_i^{-1}(\mu), \sigma_i^{-1}(\sigma_{i+1}(v_a^\ell))\} \subset \sigma_i^{-1}(V(\mathcal{K}^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)),$$

implying σ_i violates Proposition 4.6(4); $\mu \neq \sigma_{i+1}(v_a^\ell)$ since $\sigma_{i+1}^{-1}(\mu) \notin \{v_a^\ell, v_b^\ell\}$. It follows quickly that σ_{i+1} satisfies Lemma 4.7(2). This completes the induction. \square

We are now ready to prove Proposition 4.6.

Proof of Proposition 4.6. Assume (towards a contradiction) that the proposition is false, so there exists a swap sequence $\Sigma = \{\sigma_i\}_{i=0}^\lambda$ with $\sigma_0 = \sigma_s$ of minimal length λ containing a vertex that violates Proposition 4.6. It is apparent from the preceding comment that $\lambda \geq 1$. We also observe that all terms $\sigma_i \in \Sigma$ satisfy Proposition 4.4, and that σ_λ must violate at least one of the four properties of Proposition 4.6. We break into cases based on the property that the configuration σ_λ violates, and reach a contradiction in every case to deduce that none of these four properties can be violated by σ_λ . This will produce the desired contradiction on our initial assumption.

Case 1: We have $\sigma_\lambda^{-1}(\kappa_a^1) \notin V(X_a^1)$ or $\sigma_\lambda^{-1}(\kappa_b^1) \notin V(X_b^1)$. Assume that this statement holds. We only study the setting in which $\sigma_\lambda^{-1}(\kappa_a^1) \notin V(X_a^1)$; raising a contradiction when $\sigma_\lambda^{-1}(\kappa_b^1) \notin V(X_b^1)$ is analogous. To reach σ_λ from $\sigma_{\lambda-1}$, we must have $\sigma_{\lambda-1}^{-1}(\kappa_a^1) \in \{v^1, v_a^{1,2}\}$ (in particular, we must have $\lambda \geq 2$, since $\sigma_s^{-1}(\kappa_a^1) \notin \{v^1, v_a^{1,2}\}$). We break into subcases based on the value of $\sigma_{\lambda-1}^{-1}(\kappa_a^1)$.

Subcase 1.1: We have $\sigma_{\lambda-1}^{-1}(\kappa_a^1) = v^1$. Here, $\sigma_\lambda^{-1}(\kappa_a^1) \in N_{X_L}(v^1) \cap V(\mathcal{C}_b^1)$. Recall that

$$N_{Y_L}(\kappa_a^1) = (V(\mathcal{S}_a^1) \setminus \{\kappa_a^1\}) \cup V(\mathcal{K}^1).$$

Since $\sigma_{\lambda-1}$ satisfies Proposition 4.4 (on \mathcal{C}_b^1), the vertex $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^1))$ that κ_a^1 swaps with to reach σ_λ from $\sigma_{\lambda-1}$ lies in $V(\mathcal{K}^1)$. Since $\sigma_{\lambda-1}$ satisfies Proposition 4.4 (on \mathcal{C}_b^1), it must be that

$$\{\kappa_a^1, \sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^1))\} = \sigma_{\lambda-1}(V(\mathcal{C}_b^1)) \setminus (V(\mathcal{S}_b^1) \setminus \{\kappa_b^1\}).$$

Combining this with $\sigma_{\lambda-1}^{-1}(\kappa_b^1) \in V(X_b^1)$ (due to $\sigma_{\lambda-1}$ satisfying Proposition 4.6(1)), we deduce that

$$\sigma_{\lambda-1}^{-1}(\kappa_b^1) \in V(\mathcal{P}_b^1) \setminus \{v_b^1\}.$$

However, by applying Propositions 4.4 and 4.6(1-3) to $\sigma_{\lambda-1}$, and recalling our assumption that $\sigma_{\lambda-1}^{-1}(\kappa_a^1) = v^1$, we deduce that

$$\sigma_{\lambda-1}(V(\mathcal{P}_b^1) \setminus \{\kappa_b^1\}) \subset (V(\mathcal{S}_b^1) \setminus \{\kappa_b^1\}) \cup V(\mathcal{K}^1) = N_{Y_L}(\kappa_b^1),$$

so from $\sigma_{\lambda-1}$, we can swap κ_b^1 along $V(\mathcal{P}_b^1)$ onto v_b^1 , yielding a configuration $\tau \in V(\mathcal{C})$ satisfying

$$|\tau(V(\mathcal{C}_b^1)) \setminus (V(\mathcal{S}_b^1) \setminus \{\kappa_b^1\})| \geq 3,$$

contradicting Proposition 4.4. See Figure 4.6 for an illustration. In particular, this argument (with the analogue for the setting where $\sigma_{\lambda-1}^{-1}(\kappa_b^1) \notin V(X_b^1)$) concludes the study of the first three cases for $L = 1$.




Figure 4.6: Configurations in Σ used to raise a contradiction for Subcase 1.1, where we let $\mu = \sigma_{\lambda-1}(\sigma_{\lambda-1}^{-1}(\kappa_a^1))$. From $\sigma_{\lambda-1}$, swapping κ_b^1 left onto v_b^1 yields a configuration τ which violates Proposition 4.4 on \mathcal{C}_b^1 .

Subcase 1.2: We have $\sigma_{\lambda-1}^{-1}(\kappa_a^1) = v_a^{1,2}$. This subcase only applies for $L \geq 2$. Observing that we must have

$$\sigma_{\lambda-1}^{-1}(\kappa_a^1) \in N_{X_L}(v_a^{1,2}) \cap V(\mathcal{C}_a^2),$$

studying $N_{Y_L}(\kappa_a^1)$ yields $\sigma_{\lambda-1}(v_a^{1,2}) \in \{\kappa_a^2, \kappa_b^2\}$, since

$$\sigma_{\lambda-1}(v_a^{1,2}) \in V(\mathcal{S}_a^1) \setminus \{\kappa_a^1\} \text{ and } \sigma_{\lambda-1}(v_a^{1,2}) \in V(\mathcal{K}^1) \setminus \{\kappa_a^2, \kappa_b^2\}$$

imply $\sigma_{\lambda-1}$ violates Proposition 4.4 and Proposition 4.6(3), respectively. Since σ_j satisfies Propositions 4.4 and 4.6(2,3) for all $0 \leq j \leq \lambda - 1$, a case check on the types of vertices in $V(Y_L)$ and considering which of them can be in $\sigma_j(V(X^1) \setminus \{v_a^{1,2}, v_b^{1,2}\})$ implies

$$\sigma_j(V(X^1) \setminus \{v_a^{1,2}, v_b^{1,2}\}) \subset \sigma_s(V(X^1)) \text{ for all } 0 \leq j \leq \lambda - 1. \tag{4.7}$$

The observations $|\sigma_{\lambda-1}(V(X^1))| = |\sigma_s(V(X^1))|$ and $\sigma_{\lambda-1}^{-1}(\{\kappa_a^2, \kappa_b^2\}) \subset V(X^1)$ together imply that, since κ_a^1 swaps with either κ_a^2 or κ_b^2 into $N_{X_L}(v_a^{1,2}) \cap V(\mathcal{C}_a^2)$ to reach $\sigma_{\lambda-1}$ from $\sigma_{\lambda-1}$,

$$\sigma_{\lambda-1}(V(X^1)) \setminus \sigma_s(V(X^1)) \neq \emptyset, \tag{4.8}$$

while (4.7) applied to $j = \lambda - 1$ and $\sigma_{\lambda-1}(v_a^{1,2}) = \kappa_a^1 \in \sigma_s(V(X^1))$ together imply that

$$|\sigma_{\lambda-1}(V(X^1)) \setminus \sigma_s(V(X^1))| = |\sigma_s(V(X^1)) \setminus \sigma_{\lambda-1}(V(X^1))| \leq 1. \tag{4.9}$$

If it were true that $\sigma_{\lambda-1}(v_b^{1,2}) \in \sigma_s(V(X^1))$, recalling that $\sigma_{\lambda-1}(v_a^{1,2}) = \kappa_a^1 \in \sigma_s(V(X^1))$, we get

$$\sigma_{\lambda-1}(\{v_a^{1,2}, v_b^{1,2}\}) \subset \sigma_s(V(X^1)) \implies \sigma_{\lambda-1}(V(X^1) \setminus \{v_a^{1,2}, v_b^{1,2}\}) \not\subset \sigma_s(V(X^1)),$$

with the implication due to (4.8), contradicting (4.7) on $j = \lambda - 1$. Therefore,

$$\sigma_{\lambda-1}(v_b^{1,2}) \notin \sigma_s(V(X^1)).$$

This result, alongside a case check on the possible values of $\sigma_{\lambda-1}(v_b^{1,2})$ (applying Propositions 4.4 and 4.6(2,3) to $\sigma_{\lambda-1}$), gives

$$\sigma_{\lambda-1}(v_b^{1,2}) \in (V(\mathcal{S}_b^2) \setminus \{\kappa_b^2\}) \cup V(\mathcal{K}^2). \quad (4.10)$$

Let σ_ξ be the final term of Σ before $\sigma_{\lambda-1}$ satisfying

$$\sigma_\xi(v_b^{1,2}) \neq \sigma_{\lambda-1}(v_b^{1,2});$$

$\xi < \lambda - 1$ is well-defined since $\sigma_{\lambda-1}(v_b^{1,2}) \notin \sigma_s(V(X^1))$. To reach $\sigma_{\xi+1}$ from σ_ξ , we swap $\sigma_{\lambda-1}(v_b^{1,2})$ with $\sigma_\xi(v_b^{1,2})$, where

$$\sigma_\xi^{-1}(\sigma_{\lambda-1}(v_b^{1,2})) \in N_{X_L}(v_b^{1,2}) \cap V(\mathcal{C}_b^2).$$

Indeed, see (4.10); if we had that

$$\sigma_\xi^{-1}(\sigma_{\lambda-1}(v_b^{1,2})) \in N_{X_L}(v_b^{1,2}) \cap V(\mathcal{C}_b^1),$$

σ_ξ would violate Proposition 4.4 on \mathcal{C}_b^2 if $\sigma_{\lambda-1}(v_b^{1,2}) \in V(\mathcal{S}_b^2) \setminus \{\kappa_b^2\}$ and Proposition 4.6(2,3) if $\sigma_{\lambda-1}(v_b^{1,2}) \in V(\mathcal{K}^2)$. By the definition of ξ , $\sigma_j(v_b^{1,2})$ remains unchanged for $\xi + 1 \leq j \leq \lambda - 1$. Furthermore, from (4.10), we observe that

$$\sigma_\xi(v_b^{1,2}) \in \{\kappa_a^2, \kappa_b^2\} \cup V(\mathcal{K}^2), \quad (4.11)$$

since the statements

$$\sigma_\xi(v_b^{1,2}) \in (V(\mathcal{S}_a^3) \setminus \{\kappa_a^3\}) \cup (V(\mathcal{S}_b^3) \setminus \{\kappa_b^3\}) \text{ and } \sigma_\xi(v_b^{1,2}) \in V(\mathcal{K}^3)$$

would result in σ_ξ violating Propositions 4.4 and 4.6(2,3), respectively. If $\sigma_\xi(v_b^{1,2}) \in \{\kappa_a^2, \kappa_b^2\}$, then

$$\sigma_\lambda^{-1}(\sigma_\xi(v_b^{1,2})) \in V(X^1);$$

this is immediate if $\sigma_\xi(v_b^{1,2}) = \sigma_\lambda(v_a^{1,2})$ (recall that $\sigma_\lambda(v_a^{1,2}) \in \{\kappa_a^2, \kappa_b^2\}$), and if $\sigma_\xi(v_b^{1,2}) \neq \sigma_\lambda(v_a^{1,2})$, the assumption $\sigma_{\lambda-1}^{-1}(\sigma_\xi(v_b^{1,2})) = \sigma_\lambda^{-1}(\sigma_\xi(v_b^{1,2})) \notin V(X^1)$ (we swap $\sigma_\lambda(v_a^{1,2})$ and κ_a^1 to reach σ_λ from $\sigma_{\lambda-1}$), alongside $\sigma_{\lambda-1}^{-1}(\sigma_\lambda(v_a^{1,2})) \in N_{X_L}(v_a^{1,2}) \cap V(\mathcal{C}_a^2)$, would contradict (4.9), since we would have

$$\{\sigma_\xi(v_b^{1,2}), \sigma_\lambda(v_a^{1,2})\} \subseteq \sigma_s(V(X^1)) \setminus \sigma_{\lambda-1}(V(X^1)).$$

Thus, $\sigma_\xi(v_b^{1,2})$ traverses a path from $v_b^{1,2}$ to $v_a^{1,2}$, not involving $v_b^{1,2}$ past σ_ξ , as we go from σ_ξ to σ_λ . Certainly, this traversal swaps $\sigma_\xi(v_b^{1,2})$ along both $V(\mathcal{C}_a^2)$ and $V(\mathcal{C}_b^2)$. Suppose $\sigma_\xi(v_b^{1,2}) = \kappa_a^2$.

Due to (4.10), $\sigma_\xi(v_b^{1,2}) = \kappa_a^2$ must have swapped with $\sigma_{\lambda-1}(v_b^{1,2}) \in V(\mathcal{K}^2)$ to reach $\sigma_{\xi+1}$ from σ_ξ . Let $\zeta > \xi + 1$ be the earliest such index satisfying

$$\sigma_\zeta(\sigma_{\xi+1}^{-1}(\kappa_a^2)) \neq \kappa_a^2;$$

ζ is well-defined since $\sigma_\xi(v_b^{1,2}) = \kappa_a^2$ swaps along both $V(\mathcal{C}_a^2)$ and $V(\mathcal{C}_b^2)$ to reach σ_λ . The vertex $\sigma_\zeta(\sigma_{\xi+1}^{-1}(\kappa_a^2))$ must have swapped with κ_a^2 to reach σ_ζ from $\sigma_{\zeta-1}$. Since $\sigma_\zeta(v_b^{1,2}) = \sigma_{\lambda-1}(v_b^{1,2}) \in V(\mathcal{K}^2)$, $\sigma_\zeta(\sigma_{\xi+1}^{-1}(\kappa_a^2)) \in N_{Y_L}(\kappa_a^2)$, and κ_a^2 are all not in $V(\mathcal{S}_b^2) \setminus \{\kappa_b^2\}$, we have

$$|\sigma_\zeta(V(\mathcal{C}_b^2)) \setminus (V(\mathcal{S}_b^2) \setminus \{\kappa_b^2\})| \geq 3,$$

contradicting Proposition 4.4. See Figure 4.7a for an illustration. Thus, $\sigma_\xi(v_b^{1,2}) = \kappa_b^2$. Let $\zeta > \xi + 1$ be the earliest such index satisfying

$$\sigma_\zeta^{-1}(\kappa_b^2) \in V(\mathcal{C}_a^2) \setminus N_{X_L}[\text{bd}(\mathcal{C}_a^2)];$$

ζ is well-defined since κ_b^2 goes from $v_b^{1,2}$ to $v_a^{1,2}$ to reach σ_λ . Here, κ_b^2 must have swapped with κ_a^2 to reach σ_ζ from $\sigma_{\zeta-1}$: as in the preceding case, κ_b^2 would be “stuck” otherwise, due to σ_ζ satisfying Proposition 4.4 (on \mathcal{C}_a^2). But then σ_ζ would violate Proposition 4.6(4) on $\ell = 1$, namely since $\{\kappa_a^2, \kappa_b^2\} \subset V(\mathcal{K}^1)$, which implies

$$|\sigma_\zeta^{-1}(V(\mathcal{K}^1)) \setminus (V(\mathcal{P}_a^1) \cup V(\mathcal{P}_b^1))| \geq 2.$$

See Figure 4.7b for an illustration. So, by (4.11), we must have $\sigma_\xi(v_b^{1,2}) \in V(\mathcal{K}^2)$. Since we swap $\sigma_{\lambda-1}(v_b^{1,2})$ with $\sigma_\xi(v_b^{1,2})$ to reach $\sigma_{\xi+1}$ from σ_ξ , it follows from (4.10) that $\sigma_{\lambda-1}(v_b^{1,2}) \in V(\mathcal{K}^2)$, since there is no element of $V(\mathcal{K}^2)$ (in particular, $\sigma_\xi(v_b^{1,2})$) that can swap with an element of $V(\mathcal{S}_b^2) \setminus \{\kappa_b^2\}$. But then σ_ξ violates Proposition 4.6(4) (on $\ell = 2$), which is our final contradiction in this case. We conclude that Proposition 4.6(1) cannot have been the property violated by σ_λ .

Case 2: For some $\ell \geq 2$, we have $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell) \notin V(X^{\ell-1}) \cup V(X^\ell)$ or $\sigma_{\lambda-1}^{-1}(\kappa_b^\ell) \notin V(X^{\ell-1}) \cup V(X^\ell)$. This case is relevant only for $L \geq 2$. Assume this statement holds for some $2 \leq \ell \leq L$. We only study the setting in which $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell) \notin V(X^{\ell-1}) \cup V(X^\ell)$. Raising a contradiction when $\sigma_{\lambda-1}^{-1}(\kappa_b^\ell) \notin V(X^{\ell-1}) \cup V(X^\ell)$ is entirely analogous. Notice that

$$\sigma_{\lambda-1}^{-1}(\kappa_a^\ell) \in \{v_a^{\ell-2, \ell-1}, v_b^{\ell-2, \ell-1}, v_a^{\ell, \ell+1}, v_b^{\ell, \ell+1}\}$$

(precisely, the RHS above is the subset of these vertices defined for ℓ). To reach σ_λ from $\sigma_{\lambda-1}$, the vertex $\sigma_{\lambda-1}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell))$ that κ_a^ℓ swaps with satisfies

$$\sigma_{\lambda-1}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) \in \{\kappa_a^{\ell-1}, \kappa_b^{\ell-1}\} \cup V(\mathcal{K}^\ell) \cup (V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\}), \tag{4.12}$$

as $\sigma_{\lambda-1}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) \in V(\mathcal{K}_b^{\ell-1})$ would cause $\sigma_{\lambda-1}$ to violate Proposition 4.6(4), since we then get

$$\{\kappa_a^\ell, \sigma_{\lambda-1}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell))\} \subset \sigma_{\lambda-1}^{-1}(V(\mathcal{K}^{\ell-1})) \setminus (V(\mathcal{P}_a^{\ell-1}) \cup V(\mathcal{P}_b^{\ell-1}))$$

which implies that

$$|\sigma_{\lambda-1}^{-1}(V(\mathcal{K}^{\ell-1})) \setminus (V(\mathcal{P}_a^{\ell-1}) \cup V(\mathcal{P}_b^{\ell-1}))| \geq 2.$$

We break into subcases based on the value of $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)$.

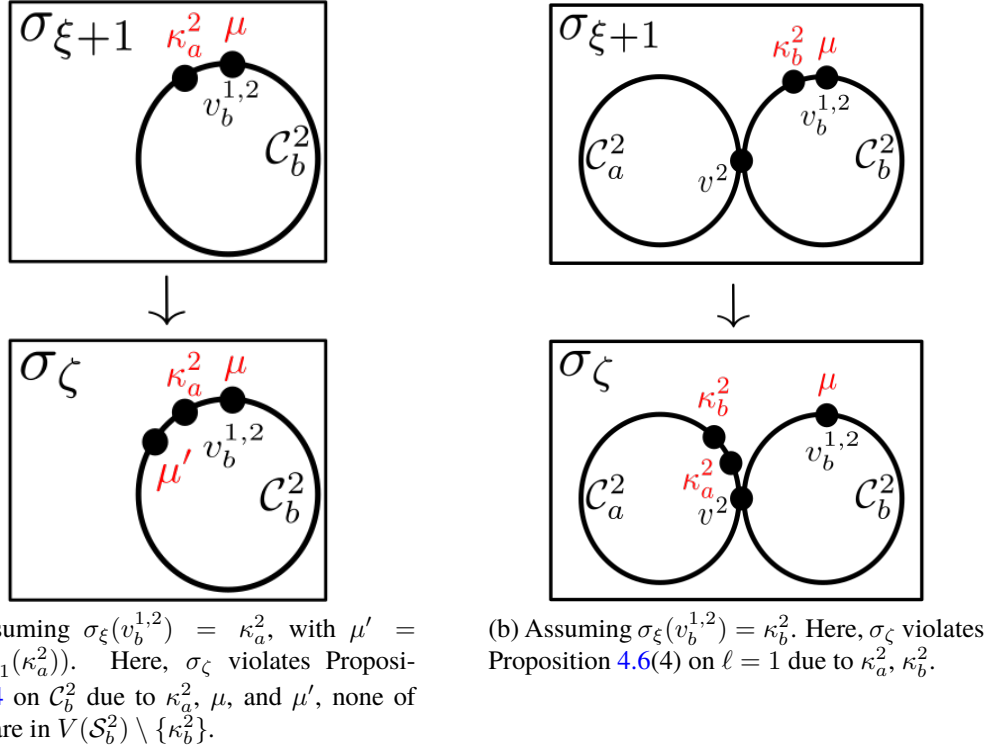


Figure 4.7: Configurations in Σ used to raise a contradiction for Subcase 1.2 when we assume that $\sigma_\xi(v_b^{1,2}) \in \{\kappa_a^2, \kappa_b^2\}$. We let $\mu = \sigma_{\lambda-1}(v_b^{1,2})$.

Subcase 2.1: We have $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell) \in \{v_a^{\ell-2, \ell-1}, v_b^{\ell-2, \ell-1}\}$. This subcase applies for $\ell \geq 3$. The vertex which κ_a^ℓ swaps onto satisfies

$$\sigma_\lambda^{-1}(\kappa_a^\ell) \in (N_{X_L}(v_a^{\ell-2, \ell-1}) \cup N_{X_L}(v_b^{\ell-2, \ell-1})) \cap V(X^{\ell-2}).$$

From (4.12), we deduce that the vertex κ_a^ℓ swaps with to reach σ_λ from $\sigma_{\lambda-1}$ satisfies

$$\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell)) \in \{\kappa_a^{\ell-1}, \kappa_b^{\ell-1}\},$$

since the statements

$$\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell)) \in V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\} \text{ and } \sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell)) \in V(\mathcal{K}^\ell)$$

respectively imply that $\sigma_{\lambda-1}$ violates Proposition 4.4 and Proposition 4.6(2,3). Proceeding backwards in Σ , $\sigma_{\lambda-2} \neq \sigma_\lambda$ ($\sigma_{\lambda-1}^{-1}(\kappa_a^\ell) \neq \sigma_s^{-1}(\kappa_a^\ell)$ implies $\sigma_{\lambda-1} \neq \sigma_s$, so $\sigma_{\lambda-2}$ is well-defined, and λ is minimal). Now, if we had that

$$\sigma_{\lambda-2}^{-1}(\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell))) = \sigma_\lambda^{-1}(\kappa_a^\ell) \text{ and } \sigma_{\lambda-2}^{-1}(\kappa_a^\ell) = \sigma_{\lambda-1}^{-1}(\kappa_a^\ell),$$

then swapping $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell))$ with κ_a^ℓ directly from $\sigma_{\lambda-2}$ would contradict λ being minimal. Thus, we have

$$\sigma_{\lambda-2}^{-1}(\kappa_a^\ell) \in N_{X_L}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) \cap V(X^{\ell-1}),$$

since $\sigma_{\lambda-2}$ satisfies Proposition 4.4 and neither $\sigma_{\lambda-1}(\sigma_{\lambda}^{-1}(\kappa_a^\ell))$ nor κ_a^ℓ can swap with vertices in

$$(V(\mathcal{S}_a^{\ell-2}) \setminus \{\kappa_a^{\ell-2}\}) \cup (V(\mathcal{S}_b^{\ell-2}) \setminus \{\kappa_b^{\ell-2}\}).$$

But any vertex in $N_{Y_L}(\kappa_a^\ell)$ with which κ_a^ℓ can swap to reach $\sigma_{\lambda-2}$ from $\sigma_{\lambda-1}$ raises a contradiction: a vertex of $\{\kappa_a^{\ell-1}, \kappa_b^{\ell-1}\} \cup V(\mathcal{K}^{\ell-1})$ implies $\sigma_{\lambda-1}$ violates Proposition 4.6(4) (respectively, on layers $\ell-2$ and $\ell-1$, due to $\sigma_{\lambda-1}(\sigma_{\lambda}^{-1}(\kappa_a^\ell))$ and κ_a^ℓ), a vertex of $V(\mathcal{K}^\ell)$ implies $\sigma_{\lambda-1}$ violates Proposition 4.6(2,3), and a vertex in $V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\}$ implies $\sigma_{\lambda-1}$ violates Proposition 4.4. See Figure 4.8 for an illustration.




Figure 4.8: Configurations in Σ used to raise a contradiction for Subcase 2.1, illustrated for $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell) = v_a^{\ell-2, \ell-1}$ and $\sigma_{\lambda-1}(\sigma_{\lambda}^{-1}(\kappa_a^\ell)) = \kappa_a^{\ell-1}$. Here, κ_a^ℓ must swap with a vertex $\mu \in N_{Y_L}(\kappa_a^\ell)$ to reach $\sigma_{\lambda-2}$ from $\sigma_{\lambda-1}$, for which all possibilities of μ raise a contradiction.

Subcase 2.2: We have $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell) \in \{v_a^{\ell, \ell+1}, v_b^{\ell, \ell+1}\}$. This subcase applies for $2 \leq \ell < L$. The vertex κ_a^ℓ swaps onto satisfies

$$\sigma_{\lambda}^{-1}(\kappa_a^\ell) \in (N_{X_L}(v_a^{\ell, \ell+1}) \cup N_{X_L}(v_b^{\ell, \ell+1})) \cap V(X^{\ell+1}).$$

From (4.12), we deduce that the vertex κ_a^ℓ swaps with to reach σ_{λ} from $\sigma_{\lambda-1}$ satisfies

$$\sigma_{\lambda-1}(\sigma_{\lambda}^{-1}(\kappa_a^\ell)) \in \{\kappa_a^{\ell+1}, \kappa_b^{\ell+1}\}, \tag{4.13}$$

since the statements

$$\sigma_{\lambda-1}(\sigma_{\lambda}^{-1}(\kappa_a^\ell)) \in V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\}, \quad \sigma_{\lambda-1}(\sigma_{\lambda}^{-1}(\kappa_a^\ell)) \in \{\kappa_a^{\ell-1}, \kappa_b^{\ell-1}\} \cup (V(\mathcal{K}^\ell) \setminus \{\kappa_a^{\ell+1}, \kappa_b^{\ell+1}\})$$

imply that $\sigma_{\lambda-1}$ violates Proposition 4.4 and Proposition 4.6(1-3), respectively. Let σ_ξ , with $\xi < \lambda - 1$, be the last term in Σ before $\sigma_{\lambda-1}$ satisfying

$$\sigma_\xi^{-1}(\sigma_{\lambda-1}(\sigma_{\lambda}^{-1}(\kappa_a^\ell))) \in V(X^\ell);$$

ξ is well-defined since (see (4.13)) $\sigma_s^{-1}(\{\kappa_a^{\ell+1}, \kappa_b^{\ell+1}\}) \subset V(X^\ell)$, which also implies $\sigma_s \neq \sigma_{\lambda-1}$ since

$$\sigma_{\lambda-1}^{-1} \left(\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell)) \right) = \sigma_\lambda^{-1}(\kappa_a^\ell) \notin V(X^\ell).$$

By the definition of ξ ,

$$\sigma_j^{-1} \left(\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell)) \right) \notin V(X^\ell) \text{ for } \xi + 1 \leq j \leq \lambda - 1. \quad (4.14)$$

Since $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell))$ traverses a path to $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)$ as we go from $\sigma_{\xi+1}$ to σ_λ , not involving $V(X^\ell)$ until σ_λ , we further deduce that

$$\sigma_\xi^{-1} \left(\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell)) \right) = \sigma_{\lambda-1}^{-1}(\kappa_a^\ell),$$

since $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell))$ cannot traverse a path from $\{v_a^{\ell,\ell+1}, v_b^{\ell,\ell+1}\} \setminus \{\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)\}$ to $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)$ as we swap from $\sigma_{\xi+1}$ to σ_λ without violating Proposition 4.4 or 4.6(2,4).⁹ Thus, from σ_ξ to $\sigma_{\xi+1}$, $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell))$ swaps into

$$N_{X_L}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) \cap V(X^{\ell+1})$$

from $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)$, and since $\sigma_{\xi+1}$ satisfies Proposition 4.6(2,3), a case check on $N_{Y_L}(\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell)))$ (see (4.13)) yields

$$\sigma_{\xi+1}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) \in (V(\mathcal{S}_a^{\ell+1}) \setminus \{\kappa_a^{\ell+1}\}) \cup (V(\mathcal{S}_b^{\ell+1}) \setminus \{\kappa_b^{\ell+1}\}) \cup V(\mathcal{K}^{\ell+1}).$$

If it were true that

$$\sigma_{\xi+1}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) \in (V(\mathcal{S}_a^{\ell+1}) \setminus \{\kappa_a^{\ell+1}\}) \cup (V(\mathcal{S}_b^{\ell+1}) \setminus \{\kappa_b^{\ell+1}\}), \quad (4.15)$$

then it must be that $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell))$ is the corresponding knob vertex. So from (4.14), we deduce that $\sigma_j(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell))$ would be fixed for $\xi + 1 \leq j \leq \lambda - 1$. Taking $j = \xi + 1$ and $j = \lambda - 1$ would imply

$$\sigma_{\xi+1}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) = \kappa_a^\ell,$$

contradicting (4.15). Therefore, $\sigma_{\xi+1}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) \in V(\mathcal{K}^{\ell+1})$. We now inductively establish that

$$\sigma_j(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) \in V(\mathcal{K}^{\ell+1}) \text{ for } \xi + 1 \leq j \leq \lambda - 1 \quad (4.16)$$

by showing that it is fixed for all such j . Assume for some j satisfying $\xi + 1 \leq j < \lambda - 1$ that σ_j satisfies this claim. Then to reach σ_{j+1} from σ_j , $\sigma_j(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell))$ cannot swap with either $\{\kappa_a^{\ell+1}, \kappa_b^{\ell+1}\}$ (see (4.13)); σ_{j+1} would violate either (4.14) or Proposition 4.6(4) for layer ℓ , depending on whether it swaps with $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell))$ or not, respectively), another vertex in $V(\mathcal{K}^{\ell+1})$ (σ_{j+1} would violate Proposition 4.6(4) for layer $\ell + 1$), or a vertex in

$$(V(\mathcal{S}_a^{\ell+2}) \setminus \{\kappa_a^{\ell+2}\}) \cup (V(\mathcal{S}_b^{\ell+2}) \setminus \{\kappa_b^{\ell+2}\}) \cup V(\mathcal{K}^{\ell+2})$$

(for the setting $\sigma_j(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) \in \{\kappa_a^{\ell+2}, \kappa_b^{\ell+2}\}$, if it applies; σ_{j+1} would violate Proposition 4.4 or Proposition 4.6(2,3)). This completes the induction. Now, (4.16) on $j = \lambda - 1$ raises a contradiction, since $\kappa_a^\ell \notin V(\mathcal{K}^{\ell+1})$. See Figure 4.9 for an illustration. This is our final contradiction in this case. We conclude that Proposition 4.6(2) cannot have been the property violated by σ_λ .

⁹This can be proved using arguments essentially identical to those in Subcase 1.2.




Figure 4.9: Configurations in Σ used to raise a contradiction for Subcase 2.2, illustrated for $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell) = v_a^{\ell,\ell+1}$ and $\sigma_{\lambda-1}(\sigma_{\lambda-1}^{-1}(\kappa_a^\ell)) = \kappa_a^{\ell+1}$. We let $\mu = \sigma_{\xi+1}(v_a^{\ell,\ell+1})$. For $\xi + 1 \leq j \leq \lambda - 1$, we have $\sigma_j(v_a^{\ell,\ell+1}) \in V(\mathcal{K}^{\ell+1})$, contradicting $\sigma_{\lambda-1}^{-1}(\kappa_a^\ell) = v_a^{\ell,\ell+1}$.

Case 3: There exists $\ell \in [L - 1]$ and $\mu \in V(\mathcal{K}^\ell) \setminus \{\kappa_a^{\ell+1}, \kappa_b^{\ell+1}\}$ such that $\sigma_\lambda^{-1}(\mu) \notin V(X^\ell)$, or there exists $\mu \in V(\mathcal{K}^L)$ such that $\sigma_\lambda^{-1}(\mu) \notin V(X^L)$. This case is relevant only for $L \geq 2$. The proceeding argument raises a contradiction both when assuming the existence of $\ell \in [L - 1]$ for which there exists $\mu \in V(\mathcal{K}^\ell) \setminus \{\kappa_a^{\ell+1}, \kappa_b^{\ell+1}\}$ such that $\sigma_\lambda^{-1}(\mu) \notin V(X^\ell)$, and also when assuming the existence of $\mu \in V(\mathcal{K}^L)$ such that $\sigma_\lambda^{-1}(\mu) \notin V(X^L)$, taking $\ell = L$.

Observe that (where, more precisely, the RHS is the subset that is well-defined for ℓ)

$$\sigma_{\lambda-1}^{-1}(\mu) \in \{v_a^{\ell-1,\ell}, v_b^{\ell-1,\ell}, v_a^{\ell,\ell+1}, v_b^{\ell,\ell+1}\},$$

and also that the vertex $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\mu))$ that μ swaps with to reach σ_λ from $\sigma_{\lambda-1}$ satisfies

$$\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\mu)) \in \{\kappa_a^\ell, \kappa_b^\ell\}, \tag{4.17}$$

since $N_{Y_L}(\mu) \subset V(\mathcal{K}^\ell) \cup \{\kappa_a^\ell, \kappa_b^\ell\}$, and $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\mu)) \in V(\mathcal{K}^\ell)$ would imply $\sigma_{\lambda-1}$ violates Proposition 4.6(4) on layer ℓ . If $\sigma_{\lambda-1}^{-1}(\mu) \in \{v_a^{\ell,\ell+1}, v_b^{\ell,\ell+1}\}$ (valid for $\ell < L$), then we would have that

$$\sigma_\lambda^{-1}(\mu) \in (N_{X_L}(v_a^{\ell,\ell+1}) \cup N_{X_L}(v_b^{\ell,\ell+1})) \cap V(X^{\ell+1}),$$

from which (4.17) implies that $\sigma_{\lambda-1}$ violates Proposition 4.6(1) if $\ell = 1$ and Proposition 4.6(2) if $\ell \geq 2$. Thus, it must be that $\ell \geq 2$ and $\sigma_{\lambda-1}^{-1}(\mu) \in \{v_a^{\ell-1,\ell}, v_b^{\ell-1,\ell}\}$, so that

$$\sigma_\lambda^{-1}(\mu) \in (N_{X_L}(v_a^{\ell-1,\ell}) \cup N_{X_L}(v_b^{\ell-1,\ell})) \cap V(X^{\ell-1}). \tag{4.18}$$

Proceeding backwards in Σ , $\sigma_{\lambda-2} \neq \sigma_\lambda(\sigma_{\lambda-1}^{-1}(\mu)) \neq \sigma_s^{-1}(\mu)$ implies that $\sigma_{\lambda-1} \neq \sigma_s$, so $\sigma_{\lambda-2}$ is well-defined, and λ is minimal). If neither μ nor $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\mu))$ were swapped to reach $\sigma_{\lambda-2}$

from $\sigma_{\lambda-1}$, swapping them directly from $\sigma_{\lambda-2}$ would contradict λ being minimal. Furthermore, from (4.17) and (4.18), μ swaps onto

$$N_{X_L}(\sigma_{\lambda-1}^{-1}(\mu)) \cap V(X^\ell)$$

to reach $\sigma_{\lambda-2}$ from $\sigma_{\lambda-1}$, since $\sigma_{\lambda-1}$ satisfies Proposition 4.4 and neither μ nor $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\kappa_a^\ell))$ can swap with vertices in the set

$$(V(\mathcal{S}_a^{\ell-1}) \setminus \{\kappa_a^{\ell-1}\}) \cup (V(\mathcal{S}_b^{\ell-1}) \setminus \{\kappa_b^{\ell-1}\}).$$

But the vertex $\sigma_{\lambda-2}(\sigma_{\lambda-1}^{-1}(\mu))$ that μ swaps with to reach $\sigma_{\lambda-2}$ from $\sigma_{\lambda-1}$ implies that $\sigma_{\lambda-2}$ violates Proposition 4.6(4) on layer $\ell - 1$ if

$$\sigma_{\lambda-2}(\sigma_{\lambda-1}^{-1}(\mu)) \in \{\kappa_a^\ell, \kappa_b^\ell\}$$

due to $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\mu))$ and $\sigma_{\lambda-2}(\sigma_{\lambda-1}^{-1}(\mu))$ (see (4.17)) and on layer ℓ if

$$\sigma_{\lambda-2}(\sigma_{\lambda-1}^{-1}(\mu)) \in V(\mathcal{K}^\ell)$$

due to μ and $\sigma_{\lambda-2}(\sigma_{\lambda-1}^{-1}(\mu))$. See Figure 4.10 for an illustration. We conclude that Proposition 4.6(3) cannot have been the property violated by σ_λ .




Figure 4.10: Configurations in Σ used to raise a contradiction in Case 3, illustrated for $\sigma_{\lambda-1}(\sigma_\lambda^{-1}(\mu)) = \kappa_a^\ell$ and $\sigma_{\lambda-1}^{-1}(\mu) = v_a^{\ell-1, \ell}$. We let $\mu' = \sigma_{\lambda-2}(\sigma_{\lambda-1}^{-1}(\mu))$. All possibilities of μ' will cause $\sigma_{\lambda-2}$ to violate Proposition 4.6(4).

Case 4: There exists $\ell \in [L]$ such that $|\sigma_\lambda^{-1}(V(\mathcal{K}^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| \geq 2$. Assume that this statement holds for some $\ell \in [L]$. We must have that

$$|\sigma_{\lambda-1}^{-1}(V(\mathcal{K}^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| = 1 \text{ and } |\sigma_\lambda^{-1}(V(\mathcal{K}^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| = 2, \quad (4.19)$$

since $|\sigma_\lambda^{-1}(V(\mathcal{K}^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| \geq 2$, λ is minimal, and for any index $1 \leq i \leq \lambda$, we have

$$|\sigma_i^{-1}(V(\mathcal{K}^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| - |\sigma_{i-1}^{-1}(V(\mathcal{K}^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| \leq 1.$$

By (4.19) and Proposition 4.6(4), there is a unique $\mu \in V(\mathcal{K}^\ell)$ such that

$$\sigma_{\lambda-1}^{-1}(\mu) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell).$$

Furthermore, there exists $\mu' \in V(\mathcal{K}^\ell)$ such that (exactly) one of the two following statements hold:

$$\sigma_{\lambda-1}(v_a^\ell) = \mu', \sigma_\lambda^{-1}(\mu') \in N_{X_L}(v_a^\ell) \setminus V(\mathcal{P}_a^\ell); \quad \sigma_{\lambda-1}(v_b^\ell) = \mu', \sigma_\lambda^{-1}(\mu') \in N_{X_L}(v_b^\ell) \setminus V(\mathcal{P}_b^\ell).$$

Studying the neighborhoods of vertices in $V(\mathcal{K}^\ell)$ yields

$$\sigma_\lambda(\sigma_{\lambda-1}^{-1}(\mu')) \in \{\kappa_a^\ell, \kappa_b^\ell\};$$

it can easily be checked that $\sigma_{\lambda-1}$ would violate one of (4.19), Proposition 4.4, or Proposition 4.6(2,3) otherwise. We will assume (the other three cases are analogous)

$$\sigma_{\lambda-1}(v_a^\ell) = \mu', \sigma_\lambda(\sigma_{\lambda-1}^{-1}(\mu')) = \kappa_a^\ell.$$

It follows from Lemma 4.7(2) that $\sigma_{\lambda-1}^{-1}(\kappa_b^\ell) \in V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}$. But if $\ell = 1$, $\sigma_{\lambda-1}$ violates Proposition 4.6(1) since

$$\sigma_{\lambda-1}^{-1}(\{\kappa_a^1, \kappa_b^1\}) \subset V(X_a^1) \setminus \{v^1\};$$

if $\ell \geq 2$, $\sigma_{\lambda-1}$ violates Proposition 4.6(4) on layer $\ell - 1$ since

$$\sigma_{\lambda-1}^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \subset \sigma_{\lambda-1}^{-1}(V(\mathcal{K}^{\ell-1})) \setminus (V(\mathcal{P}_a^{\ell-1}) \cup V(\mathcal{P}_b^{\ell-1})).$$

See Figure 4.11 for an illustration.

We conclude that Proposition 4.6(4) cannot have been the property violated by σ_λ . Together with the conclusions of the other three cases, we conclude that σ_λ satisfies all properties of Proposition 4.6, which contradicts σ_λ failing to satisfy at least one of the properties, completing the proof. \square

We can understand Propositions 4.4 and 4.6 as separating elements of $V(Y_L)$ so that for any configuration $\sigma \in V(\mathcal{C})$, specific vertices of Y_L can lie only upon specific subgraphs of X_L . In particular, for any $\ell \in [L]$, it follows from these two results that

$$\sigma(V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell) \setminus \{v_a^\ell, v_b^\ell\}) \subseteq V(\mathcal{K}^\ell) \cup \{\kappa_a^\ell, \kappa_b^\ell\}.$$

Proposition 4.6 and Lemma 4.7 together now yield the following result.




Figure 4.11: Raising a contradiction for Case 4, illustrated under the assumptions $\sigma_{\lambda-1}(v_a^\ell) = \mu'$, $\sigma_{\lambda-1}^{-1}(\mu') \in N_{X_L}(v_a^\ell) \setminus V(\mathcal{P}_a^\ell)$, and $\sigma_\lambda(\sigma_{\lambda-1}^{-1}(\mu')) = \kappa_a^\ell$. Since there exists $\mu \in V(\mathcal{K}^\ell)$ for which $\sigma_{\lambda-1}^{-1}(\mu) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$ and $\sigma_{\lambda-1}(v_a^\ell) \in V(\mathcal{K}^\ell)$, Lemma 4.7(2) yields $\sigma_{\lambda-1}^{-1}(\kappa_b^\ell) \in V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}$. This implies that $\sigma_{\lambda-1}$ violates Proposition 4.6, regardless of what the value of ℓ is.

Proposition 4.8. *For any $\sigma \in V(\mathcal{C})$, the following two statements hold.*

1. *If $\sigma^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$, then*

$$\sigma(\{v_a^\ell, v_b^\ell\}) \subset V(\mathcal{K}^\ell) \implies |\sigma^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap ((V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}) \cup (V(\mathcal{P}_b^\ell) \setminus \{v_b^\ell\}))| = 1.$$

2. *If $\sigma^{-1}(V(\mathcal{K}^\ell)) \not\subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$, then*

$$\sigma(v_a^\ell) \in V(\mathcal{K}^\ell) \implies \sigma^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap (V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}) \neq \emptyset,$$

$$\sigma(v_b^\ell) \in V(\mathcal{K}^\ell) \implies \sigma^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap (V(\mathcal{P}_b^\ell) \setminus \{v_b^\ell\}) \neq \emptyset.$$

We now prove a third invariant of any configuration in \mathcal{C} . Toward this, we begin by introducing the following notion of ordering for elements of $V(\mathcal{K}^\ell)$ in the same partite set, which is illustrated in Figure 4.12.

Definition 4.9. For $\sigma \in V(\mathcal{C})$, $\ell \in [L]$, and $\mu_1, \mu_2 \in V(\mathcal{K}^\ell)$ in the same partite set, say that μ_1 is *left* of μ_2 on σ if (exactly) one of the following holds:

1. $\sigma^{-1}(\{\mu_1, \mu_2\}) \subset V(\mathcal{P}_a^\ell)$ and $d(\sigma^{-1}(\mu_2), v_a^\ell) < d(\sigma^{-1}(\mu_1), v_a^\ell)$,
2. $\sigma^{-1}(\{\mu_1, \mu_2\}) \subset V(\mathcal{P}_b^\ell)$ and $d(\sigma^{-1}(\mu_1), v_b^\ell) < d(\sigma^{-1}(\mu_2), v_b^\ell)$,
3. $\sigma^{-1}(\mu_1) \in V(\mathcal{P}_a^\ell)$ and $\sigma^{-1}(\mu_2) \in V(\mathcal{P}_b^\ell)$.

Since $\sigma_s^{-1}(V(\mathcal{K}^\ell)) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$, it follows from Definition 4.9 that for any $\mu_1, \mu_2 \in V(\mathcal{K}^\ell)$ in the same partite set, either μ_1 is left of μ_2 on σ_s or μ_2 is left of μ_1 on σ_s . The following proposition asserts that the left relation established by σ_s cannot change for other $\sigma \in V(\mathcal{C})$.

Proposition 4.10. *Take $\ell \in [L]$ and $\mu_1, \mu_2 \in V(\mathcal{K}^\ell)$ in the same partite set, with μ_1 left of μ_2 on σ_s . If $\sigma \in V(\mathcal{C})$ is such that $\sigma^{-1}(\{\mu_1, \mu_2\}) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$, then μ_1 is left of μ_2 in σ .*




Figure 4.12: An illustration of Definition 4.9.

Proof. Let $\Sigma = \{\sigma_i\}_{i=0}^\lambda$ with $\sigma_0 = \sigma_s$ and $\sigma_\lambda = \sigma$ be a swap sequence in $\text{FS}(X_L, Y_L)$ starting from σ_s and ending at σ , where $\sigma^{-1}(\{\mu_1, \mu_2\}) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$. By Proposition 4.6(4), any $\sigma_i \in \Sigma$ satisfies

$$|\sigma_i^{-1}(V(\mathcal{K}^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| \leq 1,$$

so that in particular,

$$|\sigma_i^{-1}(\{\mu_1, \mu_2\}) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| \leq 1.$$

Consider the subsequence $\Sigma' = \{\sigma_{i_j}\}_{j=0}^{\lambda'} \subseteq \Sigma$, $\lambda' \leq \lambda$ with $i_0 = 0$ and then consisting of all configurations $\sigma_i \in \Sigma$ for which

$$|\sigma_{i_{j-1}}^{-1}(\{\mu_1, \mu_2\}) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| = 1 \text{ and } |\sigma_{i_j}^{-1}(\{\mu_1, \mu_2\}) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| = 0.$$

If μ_1 is left of μ_2 on $\sigma_{i_{\lambda'}}$, then μ_1 is left of μ_2 on σ_k for all $k \geq \lambda'$. Indeed, the construction of Σ' and $\sigma^{-1}(\{\mu_1, \mu_2\}) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$ imply that $\sigma_k^{-1}(\mu_1)$ and $\sigma_k^{-1}(\mu_2)$ remain upon the same path subgraphs in X_L for all such k , and this claim now follows if μ_1 is left of μ_2 on $\sigma_{i_{\lambda'}}$ due to Definition 4.9(3) and from $\{\mu_1, \mu_2\} \notin E(Y_L)$ otherwise. Since $\lambda \geq \lambda'$, it suffices to show that μ_1 is left of μ_2 on $\sigma_{i_{\lambda'}}$, toward which we can induct on j to show that μ_1 is left of μ_2 on σ_{i_j} for all $0 \leq j \leq \lambda'$. The statement holds for $j = 0$ by assumption, so assume μ_1 is left of μ_2 on σ_{i_j} for some $0 \leq j < \lambda'$. Take the unique vertex $\mu \in \{\mu_1, \mu_2\}$ such that

$$\sigma_{i_{j+1}-1}^{-1}(\mu) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell).$$

It is now straightforward to inductively argue, relying on the definition of Σ' , Proposition 4.6(4), and the fact that $\{\mu_1, \mu_2\} \notin E(Y_L)$, that the other vertex in $\{\mu_1, \mu_2\}$ (i.e., not μ) must remain upon the same path subgraph in X_L over all configurations σ_k for $i_j \leq k \leq i_{j+1}$. With this observation, it quickly follows, by breaking into cases based on which statement of Definition 4.9 yields μ_1 left of μ_2 on σ_{i_j} and relying on the fact that $\{\mu_1, \mu_2\} \notin E(Y_L)$, that μ_1 is left of μ_2 on $\sigma_{i_{j+1}}$. \square

We are now ready to prove the main result (in conjunction with Proposition 4.10) we will need in order to derive a lower bound on the diameter of \mathcal{C} .

Proposition 4.11. *For any configuration $\sigma \in V(\mathcal{C})$ and $\ell \in [L - 1]$,*

1. $\sigma^{-1}(\kappa_a^{\ell+1}) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell) \implies V(\mathcal{K}_b^\ell) \subset \sigma(V(\mathcal{P}_a^\ell)),$
2. $\sigma^{-1}(\kappa_b^{\ell+1}) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell) \implies V(\mathcal{K}_a^\ell) \subset \sigma(V(\mathcal{P}_a^\ell)).$

Proof. We will take $\sigma \in V(\mathcal{C})$ and $\ell \in [L - 1]$ such that $\sigma^{-1}(\kappa_a^{\ell+1}) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$; proving the latter implication when assuming $\sigma^{-1}(\kappa_b^{\ell+1}) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$ can be done analogously. By Proposition 4.6(4) and the assumption on $\sigma^{-1}(\kappa_a^{\ell+1})$,

$$|\sigma^{-1}(V(\mathcal{K}^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| = 1,$$

from which it follows that $V(\mathcal{K}_b^\ell) \subset \sigma(V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))$. To prove that $V(\mathcal{K}_b^\ell) \subset \sigma(V(\mathcal{P}_a^\ell))$, let $\Sigma = \{\sigma_i\}_{i=0}^\lambda$ be a swap sequence from $\sigma_0 = \sigma_s$ to $\sigma_\lambda = \sigma$: note that $\lambda \geq 1$, since

$$\sigma_s^{-1}(\kappa_a^{\ell+1}) \in V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell).$$

Consider the largest $\xi < \lambda$ for which

$$\sigma_\xi^{-1}(\kappa_a^{\ell+1}) \in V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell),$$

noting that $\xi < \lambda$ is well-defined, since $\sigma_s^{-1}(\kappa_a^{\ell+1}) \in V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$. It must be that

$$\sigma_\xi^{-1}(\kappa_a^{\ell+1}) \in \{v_a^\ell, v_b^\ell\}, \quad \sigma_{\xi+1}^{-1}(\kappa_a^{\ell+1}) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell). \quad (4.20)$$

Since $N_{Y_L}(\kappa_a^{\ell+1}) \cap V(\mathcal{K}_a^\ell) = \emptyset$, we deduce that

$$\sigma_\xi^{-1}(V(\mathcal{K}_a^\ell) \setminus \{\kappa_a^{\ell+1}\}) \subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell), \quad (4.21)$$

as otherwise, we would have that

$$|\sigma_{\xi+1}^{-1}(V(\mathcal{K}_a^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| \geq 2,$$

contradicting Proposition 4.6(4). From (4.20), (4.21), and Proposition 4.10, we further observe that

$$\sigma_\xi^{-1}(V(\mathcal{K}_a^\ell) \setminus \{\kappa_a^{\ell+1}\}) \subset V(\mathcal{P}_b^\ell), \quad (4.22)$$

as for any $\mu \in V(\mathcal{K}_a^\ell)$, $\kappa_a^{\ell+1}$ is left of μ on σ_ξ since $\kappa_a^{\ell+1}$ is left of μ on σ_s . By the definition of ξ ,

$$\sigma_k^{-1}(\kappa_a^{\ell+1}) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell) \text{ for } k > \xi,$$

so by Proposition 4.6(4),

$$\{\kappa_a^{\ell+1}\} = \sigma_k^{-1}(V(\mathcal{K}^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)) \text{ for } k > \xi. \quad (4.23)$$

If $V(\mathcal{K}_b^\ell) \subset \sigma_\xi(V(\mathcal{P}_a^\ell))$, (4.23) can be used to inductively prove that

$$V(\mathcal{K}_b^\ell) \subset \sigma_k(V(\mathcal{P}_a^\ell)) \text{ for } k > \xi,$$

with the induction basis following from (4.20). In particular, $V(\mathcal{K}_b^\ell) \subset \sigma(V(\mathcal{P}_a^\ell))$, which is the desired statement. Thus, we now proceed under the assumption $|V(\mathcal{K}_b^\ell) \setminus \sigma_\xi(V(\mathcal{P}_a^\ell))| \geq 1$. Further assume (towards a contradiction) that there exists

$$\mu \in \left(V(\mathcal{K}_b^\ell) \setminus \sigma_\xi(V(\mathcal{P}_a^\ell)) \right) \cap \sigma_\xi(V(\mathcal{P}_b^\ell)). \tag{4.24}$$

Then from (4.22), (4.24), and the fact that the LHS and RHS have equal cardinality,

$$\left(V(\mathcal{K}_a^\ell) \setminus \{\kappa_a^{\ell+1}\} \right) \cup \{\mu\} = \sigma_\xi(V(\mathcal{P}_b^\ell)). \tag{4.25}$$

See (4.20); (4.25) immediately raises a contradiction if $\sigma_\xi^{-1}(\kappa_a^{\ell+1}) = v_b^\ell$, and if $\sigma_\xi^{-1}(\kappa_a^{\ell+1}) = v_a^\ell$, (4.20) and Proposition 4.8(2) (the hypotheses necessary for the implication follow from (4.20) and (4.25)) imply that

$$\sigma_\xi(V(\mathcal{P}_b^\ell)) = \sigma_{\xi+1}(V(\mathcal{P}_b^\ell)) \text{ and } \sigma_{\xi+1}^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap (V(\mathcal{P}_b^\ell) \setminus \{v_b^\ell\}) \neq \emptyset,$$

respectively, raising a contradiction on (4.25). Therefore,

$$\left(V(\mathcal{K}_b^\ell) \setminus \sigma_\xi(V(\mathcal{P}_a^\ell)) \right) \cap \sigma_\xi(V(\mathcal{P}_b^\ell)) = \emptyset. \tag{4.26}$$

If it were true that $|V(\mathcal{K}_b^\ell) \setminus \sigma_\xi(V(\mathcal{P}_a^\ell))| \geq 2$, Proposition 4.6(4) would imply

$$|V(\mathcal{K}_b^\ell) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell))| \leq 1,$$

so there would exist $\mu \in V(\mathcal{K}_b^\ell) \setminus \sigma_\xi(V(\mathcal{P}_a^\ell))$ such that $\sigma_\xi^{-1}(\mu) \in V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)$, contradicting (4.26). Thus,

$$|V(\mathcal{K}_b^\ell) \setminus \sigma_\xi(V(\mathcal{P}_a^\ell))| = 1.$$

Letting μ denote the unique element in this set, it must be that $\mu \notin \sigma_\xi(V(\mathcal{P}_b^\ell))$ by (4.26), so that

$$\sigma_\xi^{-1}(\mu) \in \sigma_\xi^{-1}(V(\mathcal{K}_b^\ell)) \setminus (V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell)).$$

It thus follows from Proposition 4.6(4) that $\sigma_{\xi+1}(\sigma_\xi^{-1}(\mu)) = \mu$, so that $V(\mathcal{K}_b^\ell) \subset \sigma_{\xi+1}(V(\mathcal{P}_a^\ell))$. Now $V(\mathcal{K}_b^\ell) \subset \sigma(V(\mathcal{P}_a^\ell))$ can be established by arguing as when we assumed $V(\mathcal{K}_b^\ell) \subset \sigma_\xi(V(\mathcal{P}_a^\ell))$. \square

4.3. Extractions

Equipped with the results of Subsection 4.2, we are now ready to establish that there are two configurations in \mathcal{C} , one of them being σ_s , with distance $e^{\Omega(n)}$. The idea is to construct a series of swaps, layer-by-layer. For $\ell \in [L - 1]$, each iteration on layer $\ell + 1$ will require a certain number of iterations in layer ℓ . We formalize this in Definition 4.12 by a notion that we refer to as ℓ -extractions, which we illustrate in Figure 4.13.

Definition 4.12. For $\ell \in [L]$ and $\sigma, \tau \in V(\mathcal{C})$, we say that τ is an ℓ -extraction of σ if either:

1. $V(\mathcal{K}_a^\ell) \subset \sigma(V(\mathcal{P}_a^\ell))$ and $V(\mathcal{K}_a^\ell) \cap \tau(V(\mathcal{P}_a^\ell)) = \emptyset$, $V(\mathcal{K}_b^\ell) \subset \tau(V(\mathcal{P}_a^\ell))$,
2. $V(\mathcal{K}_b^\ell) \subset \sigma(V(\mathcal{P}_a^\ell))$ and $V(\mathcal{K}_b^\ell) \cap \tau(V(\mathcal{P}_a^\ell)) = \emptyset$, $V(\mathcal{K}_a^\ell) \subset \tau(V(\mathcal{P}_a^\ell))$.

In other words, if τ is an ℓ -extraction of σ , one of the two partite sets of $V(\mathcal{K}^\ell)$ is a subset of $\sigma(V(\mathcal{P}_a^\ell))$. Then τ “extracts” this partite set out of \mathcal{P}_a^ℓ and replaces it with the other partite set of $V(\mathcal{K}^\ell)$, which is then a subset of $\tau(V(\mathcal{P}_a^\ell))$.




Figure 4.13: An illustration of Definition 4.12. Red subgraphs/vertices corresponding to preimages of $V(\mathcal{K}_a^\ell)$, while blue subgraphs/vertices correspond to preimages of $V(\mathcal{K}_b^\ell)$. By Proposition 4.10, the relative ordering of the vertices in a partite set of $V(\mathcal{K}^\ell)$ is the same as in σ_s , so the appropriate knob vertex always lies upon the leftmost vertex in \mathcal{P}_a^ℓ .

For use in the proof of Proposition 4.14, we also introduce the following definition, corresponding to knob vertices in Y_L rotating about their corresponding cycle subgraphs in X_L . Recall from Subsection 4.1 that for all $\ell \in [L]$,

$$|V(\mathcal{S}_a^\ell)| = |V(\mathcal{S}_b^\ell)| = 15, \quad |V(\mathcal{C}_a^\ell)| = |V(\mathcal{C}_b^\ell)| = 16.$$

Definition 4.13. For $\ell \in [L]$, $\mu_a \in N_{Y_L}(\kappa_a^\ell)$ and a positive integer λ such that $\lambda \equiv 0 \pmod{16}$, a κ_a^ℓ -rotation with μ_a is a swap sequence $\{\sigma_i\}_{i=0}^\lambda$ for which $\sigma_i(V(\mathcal{C}_a^\ell)) = \{\mu_a\} \cup V(\mathcal{S}_a^\ell)$ for all $0 \leq i \leq \lambda$ and there exists an enumeration $V(\mathcal{C}_a^\ell) = \{v_0, v_1, \dots, v_{15}\}$ such that $\{v_{i-1}, v_i\} \in E(\mathcal{C}_a^\ell)$ for all $i \in [15]$ and $\sigma_j(v_i) = \kappa_a^\ell$ whenever $i \equiv j \pmod{16}$. Similarly, for $\ell \in [L]$, $\mu_b \in N_{Y_L}(\kappa_b^\ell)$ and a positive integer λ such that $\lambda \equiv 0 \pmod{16}$, a κ_b^ℓ -rotation with μ_b is a swap sequence $\{\sigma_i\}_{i=0}^\lambda$ for which $\sigma_i(V(\mathcal{C}_b^\ell)) = \{\mu_b\} \cup V(\mathcal{S}_b^\ell)$ for all $0 \leq i \leq \lambda$ and there exists an enumeration $V(\mathcal{C}_b^\ell) = \{v_0, v_1, \dots, v_{15}\}$ such that $\{v_{i-1}, v_i\} \in E(\mathcal{C}_b^\ell)$ for all $i \in [15]$ and $\sigma_j(v_i) = \kappa_b^\ell$ whenever $i \equiv j \pmod{16}$.

Note that Definition 4.13, which is illustrated in Figure 4.14, corresponds to a cyclic rotation of all elements in $\sigma_0(V(\mathcal{C}_a^\ell)) \setminus \{\kappa_a^\ell\}$ about the knob vertex κ_a^ℓ , which is fixed in the same position since $\sigma_0(v_0) = \sigma_\lambda(v_0) = \kappa_a^\ell$. The direction and length λ of this rotation depend on the enumeration of the vertices in the relevant cycle and the value $\lambda/16$, respectively.

The final configuration σ_f in \mathcal{C} for which we will argue that $d(\sigma_s, \sigma_f) = \Omega(n^{L-1})$ is going to be an L -extraction of σ_s (of the kind from Definition 4.12(1)). We begin by showing that for any $\ell \in [L]$, ℓ -extractions of σ_s exist in \mathcal{C} . This will follow as an immediate corollary of Proposition 4.14 by taking $\eta = 1$ for this value of ℓ , since σ_λ is then an ℓ -extraction of σ_s .




Figure 4.14: An illustration of a κ_a^ℓ -rotation with μ_a , where $\lambda = 3 \cdot 16 = 48$ and κ_a^ℓ rotates counterclockwise around \mathcal{C}_a^ℓ . Here, $\mu_1, \mu_2 \in V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\}$. As κ_a^ℓ rotates over \mathcal{C}_a^ℓ , it cyclically rotates all elements of $(V(\mathcal{S}_a^\ell) \setminus \{\kappa_a^\ell\}) \cup \{\mu_a\}$ about it. In this case, every such element moves three vertices clockwise along $V(\mathcal{C}_a^\ell) \setminus \{v_0\}$.

Proposition 4.14. For any positive integer η and $\ell \in [L]$, there exists a swap sequence $\{\sigma_i\}_{i=0}^\lambda$, $\sigma_0 = \sigma_s$ with a subsequence $\{\sigma_{i_j}\}_{j=0}^\eta$, $i_0 = 0$, $i_\eta = \lambda$ such that

1. for every $j \in [\eta]$, σ_{i_j} is an ℓ -extraction of $\sigma_{i_{j-1}}$;
2. for every $j \in [\eta]$ and $\mu \in V(\mathcal{K}^L)$, there exists a $\kappa_a^{L_j}$ -rotation with μ and $\kappa_b^{L_j}$ -rotation with μ that is a contiguous subsequence of $\{\sigma_i\}_{i=i_{j-1}}^{i_j}$.

Proof. We deviate from our usual practice in Subsections 4.2 and 4.3 of assuming that everything proceeds under the context of some fixed $L \geq 1$, and establish Proposition 4.14 via induction on L . Specifically, we will show by induction on $L \geq 1$ that for any fixed $L \geq 1$, Proposition 4.14 holds for the graphs X_L and Y_L . During the induction step, in another deviation from our usual practice, we shall be more explicit about the pairs of graphs and the starting configurations that we reference for sake of clarity.

We begin with the induction basis, $L = 1$. Here, $\ell = 1$ is the only value of ℓ for which Proposition 4.14 applies. Consider the following sequence of swaps from σ_s : Figure 4.15 illustrates the first three steps of this procedure.

1. Perform a κ_b^1 -rotation with $\sigma_s(v_b^1)$ to move $\sigma_s(v_b^1)$ to v^1 .
2. Perform a κ_a^1 -rotation with $\sigma_s(v_b^1)$ to move $\sigma_s(v_b^1)$ to v_a^1 .
3. Swap $\sigma_s(v_b^1)$ as far left through $V(\mathcal{P}_a^1)$ as possible, yielding a vertex $\mu \in V(\mathcal{K}_a^1)$ upon v_a^1 .
4. Perform a κ_a^1 -rotation with μ to move μ to v^1 .
5. Perform a κ_b^1 -rotation with μ to move μ to v_b^1 .
6. Swap μ as far right through $V(\mathcal{P}_b^1)$ as possible, producing a vertex in $V(\mathcal{K}_b^1)$ upon v_b^1 .

It is straightforward to conclude that repeating this algorithm 15 times (since $|V(\mathcal{K}_a^\ell)| = |V(\mathcal{K}_b^\ell)| = 15$) from σ_s (adapted to the mapping upon v_b^1 , then the mapping upon v_a^1 , for subsequent iterations) yields a 1-extraction σ_{i_1} of σ_s , namely of the kind in Definition 4.12(1), since

we have

$$\sigma_{i_1}(V(\mathcal{P}_b^1)) = \sigma_s(V(\mathcal{P}_a^1) \setminus \{v_a^1\}) = V(\mathcal{K}_a^1), \quad \sigma_{i_1}(V(\mathcal{P}_a^1) \setminus \{v_a^1\}) = \sigma_s(V(\mathcal{P}_b^1)) = V(\mathcal{K}_b^1),$$

and that for every $\mu \in V(\mathcal{K}^1)$, there exists a κ_a^1 -rotation with μ and κ_b^1 -rotation with μ that is a contiguous subsequence of the resulting swap sequence. It is similarly straightforward to see that we can repeat this algorithm to interchange the positions of $V(\mathcal{K}_a^1)$ and $V(\mathcal{K}_b^1)$ arbitrarily many times (i.e., for any positive integer η), with a κ_a^1 -rotation with μ and κ_b^1 -rotation with μ for every $\mu \in V(\mathcal{K}^1)$ executed as a contiguous subsequence of every such interchange. On even iterations of this interchange, we simply switch the roles of $V(\mathcal{K}_a^1)$ and $V(\mathcal{K}_b^1)$ in the above algorithm, resulting in 1-extractions of the kind in Definition 4.12(2).




Figure 4.15: An illustration of the first half of the sequence of swaps discussed for the induction basis, $L = 1$. Preimages of $V(\mathcal{K}_a^1)$ are colored green, and preimages of $V(\mathcal{K}_b^1)$ are colored blue. This segment of the sequence of swaps involves a κ_b^1 -rotation with $\sigma_s(v_b^1)$, a κ_a^1 -rotation with $\sigma_s(v_b^1)$, and a sequence of swaps moving $\sigma_s(v_b^1)$ left through $V(\mathcal{P}_a^1)$. A vertex $\mu_a \in V(\mathcal{K}_a^1)$ now lies upon v_a^1 ; we can similarly move μ_a to the right. Continuing until every vertex of $V(\mathcal{K}^1)$ is moved in an analogous subroutine yields a 1-extraction σ_{i_1} of σ_s . We can interchange \mathcal{K}_a^1 and \mathcal{K}_b^1 in this way arbitrarily many times.

Now assume Proposition 4.14 holds for some fixed $L \geq 1$ (i.e., for this fixed $L \geq 1$, Proposition 4.14 holds for the graphs X_L and Y_L). By the induction hypothesis applied on $\eta = 31$ and $\ell \in [L]$, we can extract a swap sequence $\{\sigma_i\}_{i=0}^\lambda$ in $V(\mathcal{C}(X_L, Y_L))$ with $\sigma_0 = \sigma_s(X_L, Y_L)$ and with a subsequence $\{\sigma_{i_j}\}_{j=0}^{31}$ satisfying Proposition 4.14. Now consider X_{L+1} and Y_{L+1} , which has corresponding starting configuration $\sigma_s(X_{L+1}, Y_{L+1})$ in the connected component $\mathcal{C}(X_{L+1}, Y_{L+1})$, which we denote σ_s and \mathcal{C} , respectively. In an abuse of notation, for the rest of the present proof we let X_L denote the first L layers of X_{L+1} and $Y_L = Y_{L+1}|_{\sigma_s(V(X_L))}$. These subgraphs are isomorphic to the graphs X_L and Y_L as they were originally defined during their construction in Subsection 4.1, and under these isomorphisms, σ_s restricted to X_L can be understood to be the same as $\sigma_s(X_L, Y_L)$ as defined in Subsection 4.1. Furthermore, the swap sequence $\{\sigma_i\}_{i=0}^\lambda$ can be understood as being in \mathcal{C} , with $\sigma_0 = \sigma_s$, if we set

$$\sigma_i(v) = \sigma_s(v) \text{ for all } v \in V(X_{L+1}) \setminus V(X_L), \quad i = 0, \dots, \lambda.$$

As such, it follows from the induction hypothesis that Proposition 4.14 holds for (X_{L+1}, Y_{L+1}) if we take $\ell \in [L]$, and all that remains is to confirm that Proposition 4.14 holds for (X_{L+1}, Y_{L+1}) for $\ell = L + 1$. In the proceeding argument, we assume that the swap sequence $\{\sigma_i\}_{i=0}^\lambda$ we extracted above using the induction hypothesis was for $\ell = L$. In a similar vein, given some $\sigma \in V(\mathcal{C})$ with $\sigma(X_L) = V(Y_L)$ and σ_i from the swap sequence $\{\sigma_i\}_{i=0}^\lambda$, define the *extension* of σ_i with respect to σ to be the configuration¹⁰ $\tau \in V(\text{FS}(X_{L+1}, Y_{L+1}))$ with

- $\tau(v) = \sigma(v)$ for all $v \in V(X_{L+1}) \setminus V(X_L)$;
- $\tau(v) = \sigma_i(v)$ for all $v \in V(X_L)$.

We will say that we *extend* σ_i with respect to σ , and will generally apply this notion en masse to subsequences of $\{\sigma_i\}_{i=0}^\lambda$ with respect to a single configuration of $V(\mathcal{C})$.

We will now construct a swap sequence $\{\sigma'_i\}_{i=0}^{\lambda'}$ in $V(\mathcal{C})$, with $\sigma'_0 = \sigma_s$, satisfying Proposition 4.14 for $\eta = 1$. This is illustrated in Figure 4.16. From the swap sequence $\{\sigma_i\}_{i=0}^\lambda$, with subsequence $\{\sigma_{i_j}\}_{j=0}^{31}$ as discussed before, consider $\{\sigma_i\}_{i=i_0}^{i_1}$, which, by the induction hypothesis, has a contiguous subsequence $\{\sigma_i\}_{i=j_1}^{k_1}$ that is a κ_b^L -rotation with κ_b^{L+1} . Let t_1 be such that $j_1 \leq t_1 \leq k_1$ and $\sigma_{t_1}(\kappa_b^{L+1}) = v_b^{L,L+1}$: the observation that such a t_1 exists follows quickly from the restrictions of Definition 4.13. We construct a swap sequence \mathcal{S}_1 in \mathcal{C} by merging $\{\sigma_i^{1,1}\}_{i=0}^{t_1-i_0}$, $\{\sigma_i^{1,2}\}_{i=0}^{z_1}$, and $\{\sigma_i^{1,3}\}_{i=0}^{i_1-t_1}$, which we now define.

1. Extend $\{\sigma_i\}_{i=i_0}^{t_1}$ with respect to σ_s , yielding $\{\sigma_i^{1,1}\}_{i=0}^{t_1-i_0}$.
2. Let $\{\sigma_i^{1,2}\}_{i=0}^{z_1}$, with $\sigma_0^{1,2} = \sigma_{t_1-i_0}^{1,1}$, be a κ_b^{L+1} -rotation with $\sigma_{t_1-i_0}^{1,1}(v_b^{L+1})$, with length such that $\sigma_{t_1-i_0}^{1,1}(v_b^{L+1})$ is moved to v^{L+1} .
3. Extend $\{\sigma_i\}_{i=t_1}^{i_1}$ with respect to $\sigma_{t_1-i_0}^{1,1}$, yielding $\{\sigma_i^{1,3}\}_{i=0}^{i_1-t_1}$.

Now take the subsequence $\{\sigma_i\}_{i=i_1}^{i_2}$ of $\{\sigma_i\}_{i=0}^\lambda$, which has contiguous subsequence $\{\sigma_i\}_{i=j_2}^{k_2}$ that is a κ_a^L -rotation with κ_a^{L+1} , and t_2 such that $j_2 \leq t_2 \leq k_2$ and $\sigma_{t_2}(\kappa_a^{L+1}) = v_a^{L,L+1}$. Construct \mathcal{S}_2 by merging $\{\sigma_i^{2,1}\}_{i=0}^{t_2-i_1}$, $\{\sigma_i^{2,2}\}_{i=0}^{t_2-i_1}$, and $\{\sigma_i^{2,3}\}_{i=0}^{t_2-i_1}$, which we now define.

1. Extend $\{\sigma_i\}_{i=i_1}^{t_2}$ with respect to $\sigma_{i_1-t_1}^{1,3}$, yielding $\{\sigma_i^{2,1}\}_{i=0}^{t_2-i_1}$.
2. Let $\{\sigma_i^{2,2}\}_{i=0}^{z_2}$, with $\sigma_0^{2,2} = \sigma_{t_2-i_1}^{2,1}$, be the result of performing a κ_a^{L+1} -rotation with $\sigma_s(v_b^{L+1})$ to move $\sigma_s(v_b^{L+1})$ to v_a^{L+1} , then swapping $\sigma_s(v_b^{L+1})$ as far left as possible across $V(\mathcal{P}_a^{L+1})$, then performing a κ_a^{L+1} -rotation to swap the resulting vertex μ upon v_a^{L+1} onto v^{L+1} .
3. Extend $\{\sigma_i\}_{i=t_2}^{i_2}$ with respect to $\sigma_{t_2-i_1}^{2,1}$, yielding $\{\sigma_i^{2,3}\}_{i=0}^{i_2-t_2}$.

It is now straightforward to see how to similarly construct the sequences $\mathcal{S}_1, \dots, \mathcal{S}_{31}$, and why we took $\eta = 31$ when appealing to the induction hypothesis: each sequence corresponding to a different vertex in $V(\mathcal{K}^{L+1})$ lying on v^{L+1} , and following \mathcal{S}_1 , we alternate the path that we “push” this vertex through. The only modification arises when we construct \mathcal{S}_{31} : during the κ_b^L -rotation with κ_b^{L+1} within $\{\sigma_i\}_{i=i_{\eta-1}}^{i_\eta}$, simply include a κ_b^{L+1} -rotation which moves the

¹⁰Whenever we construct such extensions in the forthcoming argument, it will be clear that they lie in $V(\mathcal{C})$.

vertex upon v^{L+1} onto v_b^{L+1} . Merging $\mathcal{S}_1, \dots, \mathcal{S}_{31}$ yields a sequence $\{\sigma'_i\}_{i=0}^{\lambda'}$, such that $\sigma'_{\lambda'}$ is an $(L+1)$ -extraction of $\sigma'_0 = \sigma_s$, since

$$\begin{aligned}\sigma'_{\lambda'}(V(\mathcal{P}_b^{L+1})) &= \sigma_s(V(\mathcal{P}_a^{L+1}) \setminus \{v_a^{L+1}\}) = V(\mathcal{K}_a^{L+1}), \\ \sigma'_{\lambda'}(V(\mathcal{P}_a^{L+1}) \setminus \{v_a^{L+1}\}) &= \sigma_s(V(\mathcal{P}_b^{L+1})) = V(\mathcal{K}_b^{L+1}).\end{aligned}$$

It is evident by tracing the above construction that for every $\mu \in V(\mathcal{K}^{L+1})$, there exists a κ_a^{L+1} -rotation with μ and κ_b^{L+1} -rotation with μ that is a contiguous subsequence of $\{\sigma'_i\}_{i=0}^{\lambda'}$. Thus, the swap sequence $\{\sigma'_i\}_{i=0}^{\lambda'}$ establishes that Proposition 4.14 holds for $L+1$ on $\eta = 1$.




Figure 4.16: An illustration of the sequences of swaps defined during the construction of $\{\sigma'_i\}_{i=0}^{\lambda'}$ in the induction step, on $L+1$ layers. Subgraphs/vertices corresponding to preimages of $V(\mathcal{K}_a^L)$, $V(\mathcal{K}_b^L)$, $V(\mathcal{K}_a^{L+1})$, and $V(\mathcal{K}_b^{L+1})$ are red, blue, gold, and pink, respectively. We specifically depict the construction of the sequence \mathcal{S}_2 , constructed from the subsequence $\{\sigma_i\}_{i=i_1}^{i_2}$ of the original sequence $\{\sigma_i\}_{i=0}^{\lambda}$. Initially, for $\sigma_0^{2,1}$, we have $\sigma_s(v_b^{L+1})$ upon v^{L+1} . At $\sigma_0^{2,2}$, we extend a κ_a^L -rotation with κ_a^{L+1} in $\{\sigma_i\}_{i=i_1}^{i_2}$ (which is guaranteed to exist by the induction hypothesis) so that it includes the following sequence of swaps: a κ_a^{L+1} -rotation with $\sigma_s(v_b^{L+1})$, swapping $\sigma_s(v_b^{L+1})$ left into \mathcal{P}_a^{L+1} , and a κ_a^{L+1} -rotation with the resulting μ_a on v_a^{L+1} . This will result in the configuration $\sigma_0^{2,3}$. From $\sigma_0^{2,3}$ to the final configuration in \mathcal{S}_2 , we execute the rest of $\{\sigma_i\}_{i=i_1}^{i_2}$, leading to an L -extraction of $\sigma_0^{2,1}$. Exhausting the original sequence $\{\sigma_i\}_{i=0}^{\lambda}$ by proceeding like this will yield an $(L+1)$ -extraction of σ_s , and the resulting swap sequence satisfies Proposition 4.14(2).

For general $\eta \geq 1$, we can invoke the induction hypothesis, applied to 31η , to extract a swap sequence $\{\sigma_i\}_{i=0}^{\lambda}$ in $V(\mathcal{C}(X_L, Y_L))$ with subsequence $\{\sigma_{i_j}\}_{j=0}^{31\eta}$. Then we can proceed as in the $\eta = 1$ case for every contiguous subsequence $\{\sigma_k\}_{k=31(i-1)}^{31i}$ in $\{\sigma_i\}_{i=0}^{\lambda}$, for $i \in [\eta]$, to construct a swap sequence $\{\sigma'_i\}_{i=0}^{\lambda'} \subset V(\mathcal{C})$ which establishes Proposition 4.14 for this value of η . \square

In the proof of Proposition 4.14, during the induction basis we reached a 1-extraction of σ_s by performing $\Omega(n)$ iterations of an algorithm which executed $\Omega(n^2)$ swaps.¹¹ Then in the inductive step, we reached an $(\ell + 1)$ -extraction of σ_s by taking $\Omega(n)$ ℓ -extractions of σ_s and stringing them together by appending some other swap sequences. Altogether, it follows that we found a sequence of $\Omega(n^{L+2})$ swaps to reach an L -extraction of σ_s — if this were tight, taking L to be as large as desired would be enough to answer Question 1.2 in the negative. Motivated by these ideas, we prove Proposition 4.15, which will lend itself to a lower bound on $d(\sigma_s, \sigma_f)$.

Proposition 4.15. *Fix integers $L \geq 2$ and $\ell \in [L - 1]$, and take $\sigma, \tau \in V(\mathcal{C})$ such that τ is an $(\ell + 1)$ -extraction of σ . Any swap sequence $\{\sigma_i\}_{i=0}^\lambda$ with $\sigma_0 = \sigma$ and $\sigma_\lambda = \tau$ must have a subsequence $\{\sigma_{i_j}\}_{j=0}^{25}$ such that, for $j \in [25]$, there exists a configuration $\tilde{\sigma} \in \{\sigma_i\}_{i=i_{j-1}}^{i_j}$ that is an ℓ -extraction of $\sigma_{i_{j-1}}$.*

Proof. Assume τ is an $(\ell + 1)$ -extraction of σ of the kind of Definition 4.12(1). Proposition 4.15 can be proved in the setting where τ is an $(\ell + 1)$ -extraction of σ of the kind of Definition 4.12(2) entirely analogously, where we switch the roles of several expressions corresponding to the “left and right sides” of the subgraphs $X^{\ell+1}$ and $\mathcal{K}^{\ell+1}$ in this case.

We will say that $X_a^{\ell+1}$ is the *initial subgraph* of any vertex $\mu \in V(\mathcal{K}_a^{\ell+1}) \setminus \{\kappa_a^{\ell+2}\}$ ($\mu \in V(\mathcal{K}_a^L)$ if $\ell = L - 1$) with $\sigma^{-1}(\mu) \in V(X_a^{\ell+1})$. By Proposition 4.6(3), μ leaving the initial subgraph corresponds to an (X_L, Y_L) -friendly swap where μ is upon $v^{\ell+1}$ and swaps onto some vertex in $N_{X_L}(v^{\ell+1}) \cap V(X_b^{\ell+1})$. Similarly, $X_b^{\ell+1}$ is the *initial subgraph* for $\mu \in V(\mathcal{K}_b^{\ell+1}) \setminus \{\kappa_b^{\ell+2}\}$ ($\mu \in V(\mathcal{K}_b^L)$ if $\ell = L - 1$) with $\sigma^{-1}(\mu) \in V(X_b^{\ell+1})$. By Proposition 4.6(3), μ leaving the initial subgraph corresponds to an (X_L, Y_L) -friendly swap where μ is upon $v^{\ell+1}$ and swaps onto some vertex in $N_{X_L}(v^{\ell+1}) \cap V(X_a^{\ell+1})$.

Let $\Sigma = \{\sigma_i\}_{i=0}^\lambda$ be a swap sequence with $\sigma_0 = \sigma$ and $\sigma_\lambda = \tau$. It is straightforward to show from Proposition 4.6(4) and Definition 4.12(1) that at least 26 vertices in $V(\mathcal{K}_a^{\ell+1}) \setminus \{\kappa_a^{\ell+2}, \kappa_b^{\ell+2}\}$ ($V(\mathcal{K}_a^L)$ for $\ell = L - 1$) switch to the “opposite” layer $\ell + 1$ subgraph in X_L over the course of Σ . Take any 26 such vertices $\{\mu_1, \dots, \mu_{26}\}$, indexed in the order that they first leave their initial subgraph during Σ (it is clear that at most one such vertex can leave their initial subgraph over a given swap). Construct a subsequence $\{\sigma_{i_j}\}_{j=1}^{26}$ of Σ such that, for every $j \in [26]$, i_j is the smallest index for which

- $\sigma_{i_j}^{-1}(\mu_j) = v^{\ell+1}$;
- $\sigma_{i_{j+1}}^{-1}(\mu_j)$ is not a vertex in the initial subgraph of μ_j .

Consider any $j \in [26]$ for which $\mu_j \in V(\mathcal{K}_a^{\ell+1})$. The neighborhood of μ_j is

$$N_{Y_L}(\mu_j) = V(\mathcal{K}_b^{\ell+1}) \cup \{\kappa_a^{\ell+1}, \kappa_b^{\ell+1}\}.$$

The vertex $\sigma_{i_{j+1}}(v^{\ell+1})$ that μ_j swaps with to reach $\sigma_{i_{j+1}}$ from σ_{i_j} satisfies

$$\sigma_{i_{j+1}}(v^{\ell+1}) \in \{\kappa_a^{\ell+1}, \kappa_b^{\ell+1}\},$$

¹¹Note that we inducted on L in the proof of Proposition 4.14, so the sequence of swaps we found for smaller values of L would be executed on subgraphs of (X_L, Y_L) for larger values of L . However, it is easy to verify, by tracing the construction in Subsection 4.1, that the asymptotic statements here hold regardless of the fixed value of L that we choose.

since σ_{i_j} would violate Proposition 4.6(4) (on layer $\ell + 1$) if we had that $\sigma_{i_{j+1}}(v^{\ell+1}) \in V(\mathcal{K}_b^{\ell+1})$. Assume (towards a contradiction) that $\sigma_{i_{j+1}}(v^{\ell+1}) = \kappa_a^{\ell+1}$, and let $1 \leq \xi \leq i_j$ (the lower bound is since $\sigma_s^{-1}(\mu_j) \neq v^{\ell+1}$) be the smallest such index satisfying

$$\sigma_\xi^{-1}(\mu_j) = \sigma_{i_j}^{-1}(\mu_j) = v^{\ell+1} \text{ and } \sigma_\xi^{-1}(\kappa_a^{\ell+1}) = \sigma_{i_j}^{-1}(\kappa_a^{\ell+1}) \in N_{X_L}(v^\ell) \cap V(X_b^{\ell+1}). \quad (4.27)$$

Exactly one of the two statements

- $\sigma_{\xi-1}^{-1}(\kappa_a^{\ell+1}) \neq \sigma_\xi^{-1}(\kappa_a^{\ell+1})$;
- $\sigma_{\xi-1}^{-1}(\mu_j) \neq \sigma_\xi^{-1}(\mu_j)$

is true; both being false would contradict ξ being the smallest possible, while both being true would contradict i_j being the smallest possible. But $\sigma_{\xi-1}^{-1}(\kappa_a^{\ell+1}) \neq \sigma_\xi^{-1}(\kappa_a^{\ell+1})$ would imply that $\sigma_{\xi-1}$ violates Proposition 4.4 (on $\mathcal{C}_b^{\ell+1}$), and $\sigma_{\xi-1}^{-1}(\mu_j) \neq \sigma_\xi^{-1}(\mu_j)$ would imply that $\sigma_{\xi-1}$ violates Proposition 4.6(4) (on layer ℓ if it swaps with $\kappa_b^{\ell+1}$, and on layer $\ell + 1$ if it swaps with a vertex in $V(\mathcal{K}_b^{\ell+1})$). So for all $j \in [26]$,

$$\mu_j \in V(\mathcal{K}_a^{\ell+1}) \implies \sigma_{i_{j+1}}(v^{\ell+1}) = \kappa_b^{\ell+1} \text{ and } \mu_j \in V(\mathcal{K}_b^{\ell+1}) \implies \sigma_{i_{j+1}}(v^{\ell+1}) = \kappa_a^{\ell+1}, \quad (4.28)$$

where the latter claim can be deduced from an entirely analogous argument. See Figure 4.17 for an illustration.




Figure 4.17: The two possibilities for the configuration σ_{i_j} for any $j \in [26]$. Subgraphs/vertices corresponding to preimages of $V(\mathcal{K}_a^\ell)$ and $V(\mathcal{K}_b^\ell)$ are colored red and blue, respectively. The coloring of \mathcal{P}_a^ℓ in both cases follows from Proposition 4.11.

Now consider $2 \leq j \leq 26$ for which $\mu_j \in V(\mathcal{K}_a^{\ell+1})$. For such values of j which have that $\mu_j \in V(\mathcal{K}_b^{\ell+1})$, establishing the existence of a configuration $\tilde{\sigma} \in \{\sigma_i\}_{i=i_{j-1}}^{i_j}$ that is an ℓ -extraction of $\sigma_{i_{j-1}}$ can be done entirely analogously. By (4.28), $\sigma_{i_{j+1}}(v^{\ell+1}) = \kappa_b^{\ell+1}$, so certainly

$$\sigma_{i_j}^{-1}(\kappa_b^{\ell+1}) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell),$$

and since $\kappa_b^{\ell+1} \in V(\mathcal{K}^\ell)$,

$$\sigma_{i_j}^{-1}(V(\mathcal{K}^\ell)) \not\subset V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell).$$

By Proposition 4.11(2),

$$V(\mathcal{K}_a^\ell) \subset \sigma_{i_j}(V(\mathcal{P}_a^\ell)). \tag{4.29}$$

If it were true that $V(\mathcal{K}_b^\ell) \cap \sigma_{i_j}(V(\mathcal{P}_a^\ell)) \neq \emptyset$, there would exist $\eta_1 \in V(\mathcal{K}_b^\ell)$ satisfying $\eta_1 \in \sigma_{i_j}(V(\mathcal{P}_a^\ell))$. Combined with (4.29), we would have

$$V(\mathcal{K}_a^\ell) \cup \{\eta_1\} = \sigma_{i_j}(V(\mathcal{P}_a^\ell)), \tag{4.30}$$

since the LHS is a subset of the RHS and their cardinalities are equal. In particular, it must be that $\sigma_{i_j}(v_a^\ell) \in V(\mathcal{K}^\ell)$, so Proposition 4.8(2) would imply that

$$\sigma_{i_j}^{-1}(\{\kappa_a^\ell, \kappa_b^\ell\}) \cap (V(\mathcal{P}_a^\ell) \setminus \{v_a^\ell\}) \neq \emptyset,$$

so there exists $\eta_2 \in \{\kappa_a^\ell, \kappa_b^\ell\}$ that is in $\sigma_{i_j}(V(\mathcal{P}_a^\ell))$, contradicting (4.30). So it must be that $\sigma_{i_j}^{-1}(V(\mathcal{K}_b^\ell)) \cap V(\mathcal{P}_a^\ell) = \emptyset$, i.e., that

$$V(\mathcal{K}_b^\ell) \cap \sigma_{i_j}(V(\mathcal{P}_a^\ell)) = \emptyset. \tag{4.31}$$

If $\mu_{j-1} \in V(\mathcal{K}_b^{\ell+1})$, then (4.28) implies $\sigma_{i_{j-1}+1}(v^{\ell+1}) = \kappa_a^{\ell+1}$, and so

$$\sigma_{i_{j-1}}^{-1}(\kappa_a^{\ell+1}) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell).$$

Proposition 4.11(1) thus implies

$$V(\mathcal{K}_b^\ell) \subset \sigma_{i_{j-1}}(V(\mathcal{P}_a^\ell)).$$

This statement, with (4.29) and (4.31), implies that σ_{i_j} is an ℓ -extraction of $\sigma_{i_{j-1}}$, namely of the Definition 4.12(2) kind. If $\mu_{j-1} \in V(\mathcal{K}_a^{\ell+1})$, then (4.28) implies $\sigma_{i_{j-1}+1}(v^{\ell+1}) = \kappa_b^{\ell+1}$, so

$$\sigma_{i_{j-1}}^{-1}(\kappa_b^{\ell+1}) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell).$$

Proposition 4.11(2) now implies that

$$V(\mathcal{K}_a^\ell) \subset \sigma_{i_{j-1}}(V(\mathcal{P}_a^\ell)).$$

Since i_j is the smallest possible, $\sigma_{i_{j-1}}^{-1}(\mu_j) \in V(X_a^{\ell+1})$, and it follows that $\sigma_{i_{j-1}}^{-1}(\mu_j) \in V(\mathcal{P}_a^{\ell+1})$, as $\sigma_{i_{j-1}}$ would otherwise violate Proposition 4.6(4) on layer ℓ (due to $\kappa_b^{\ell+1}$ and μ_j). It is straightforward to confirm, appealing to Proposition 4.4 on $\mathcal{C}_a^{\ell+1}$, that μ_j moves to $v^{\ell+1}$ during $\{\sigma_i\}_{i=i_{j-1}}^{i_j}$, and swaps with $\kappa_a^{\ell+1}$ upon $V(\mathcal{C}_a^{\ell+1})$ at some point in this swap sequence.¹² Thus, there exists a configuration $\tilde{\sigma} \in \{\sigma_i\}_{i=i_{j-1}}^{i_j}$ for which

$$\tilde{\sigma}^{-1}(\{\mu_j, \kappa_a^{\ell+1}\}) \subset V(\mathcal{C}_a^{\ell+1}) \text{ and } \tilde{\sigma}^{-1}(\kappa_a^{\ell+1}) \neq v_a^\ell,$$

¹²This can be proved using ideas and arguments which are essentially identical to those that were carried out in Subcase 1.2 of the proof of Proposition 4.6.

from which it immediately follows that

$$\tilde{\sigma}^{-1}(\kappa_a^{\ell+1}) \notin V(\mathcal{P}_a^\ell) \cup V(\mathcal{P}_b^\ell),$$

and Proposition 4.11(1) implies

$$V(\mathcal{K}_b^\ell) \subset \tilde{\sigma}(V(\mathcal{P}_a^\ell)).$$

This, with (4.29) and (4.31), implies that $\tilde{\sigma}$ is an ℓ -extraction of $\sigma_{i_{j-1}}$, namely of the Definition 4.12(2) kind.

Therefore, taking $\{\sigma_{i_j}\}_{j=1}^{26}$ yields the desired subsequence of $\{\sigma_i\}_{i=0}^\lambda$. \square

4.4. Proof of Theorem 1.3

We finally derive the desired lower bound on the diameter of \mathcal{C} .

Theorem 1.3. *For all $n \geq 2$, there exist n -vertex graphs X and Y such that $\text{FS}(X, Y)$ has a connected component with diameter $e^{\Omega(n)}$.*

Proof. For $L \geq 2$, take X_L, Y_L on $58L + 2$ vertices (see Subsection 4.1). For $\ell \in [L - 1]$, define

$$\lambda_{(L,n)}(\ell) := \min \{d(\sigma, \tau) : \sigma, \tau \in V(\mathcal{C}), \tau \text{ is an } \ell\text{-extraction of } \sigma\}. \quad (4.32)$$

It follows from Proposition 4.15 that for all $\ell \in [L - 1]$,

$$\lambda_{(L,n)}(\ell + 1) \geq 25\lambda_{(L,n)}(\ell). \quad (4.33)$$

Let $\sigma_f \in V(\mathcal{C})$ be such that σ_f is an L -extraction of σ_s , which exists by Proposition 4.14. By (4.32) and (4.33),

$$d(\sigma_s, \sigma_f) \geq \lambda_{(L,n)}(L) \geq 25\lambda_{(L,n)}(L - 1) \geq \dots \geq 25^{L-1}\lambda_{(L,n)}(1) \geq 25^{L-1}.$$

Now, for $n \geq 60$, fix $L = \lfloor (n - 2)/58 \rfloor$ (here, $L \geq 1$), and construct n -vertex graphs \tilde{X}_n, \tilde{Y}_n by adding $n' = n - (58L + 2)$ isolated vertices to X_L and Y_L , respectively. Let $\mathcal{C}(\tilde{X}_n, \tilde{Y}_n)$ denote the connected component of $\text{FS}(\tilde{X}_n, \tilde{Y}_n)$ containing the configuration resulting from placing $V(\tilde{Y}_n)$ upon $V(\tilde{X}_n)$ as usual (i.e., under the starting configuration as defined in Subsection 4.1), and then placing the $n - n'$ isolated vertices in \tilde{Y}_n upon the $n - n'$ isolated vertices of \tilde{X}_n in some way. It easily follows from our construction that $58L + 2 \leq n \leq 58L + 58$, so

$$d(\sigma_s, \sigma_f) \geq 25^{L-1} = e^{\Omega(n)}.$$

By accounting for the values $2 \leq n \leq 59$, which may weaken the constant implicit in the $\Omega(n)$ term, the desired result now follows immediately. \square

To conclude Section 4, we mention an especially notable implication of Theorem 1.3 in the study of random walks on friends-and-strangers graphs. In the proceeding discussion, to avoid distracting from the nature of this article, we elect to be terse and do not define many of the objects we consider; we refer the reader to [Lov93] for a classical survey of random walks on graphs. We begin by providing the following Definition 4.16. The fact that there is a natural discrete-time Markov chain associated to a friends-and-strangers graph was observed in passing in [ADK23, Section 7], and an investigation of its mixing properties was separately proposed in [Alo21].

Definition 4.16. Let X and Y be n -vertex graphs. The *friends-and-strangers Markov chain* of X and Y is the discrete-time Markov chain whose state space is $V(\text{FS}(X, Y))$ and such that at each time step, a pair of friends standing upon adjacent vertices is chosen uniformly at random amongst all such pairs and swap places with probability $1/2$.

The friends-and-strangers Markov chain of X and Y , which models a lazy random walk on a connected component of $\text{FS}(X, Y)$, is aperiodic by construction (recall from Proposition 2.3(2) that friends-and-strangers graphs are bipartite, which warrants the laziness condition in Definition 4.16 if we would like to discuss mixing to stationarity) and irreducible when restricted to a connected component of $\text{FS}(X, Y)$. From a different perspective, Definition 4.16 may be interpreted as the generalization of a natural discrete-time variant of the interchange process¹³ (sometimes called the random stirring process) where we include the condition that arbitrary pairs of particles may be forbidden from swapping positions. The friends-and-strangers Markov chain has received substantial attention under certain restricted settings (chiefly that in which $Y = K_n$, and oftentimes under the context of card shuffling Markov chains), in which classical polynomial upper bounds on the mixing time (in total variation distance) of the underlying Markov chain are known; see [Ald83, AD86, DS81, DSC93, Jon12, Mat88, Wil04]. An immediate corollary of Theorem 1.3, which might be thought of as its natural stochastic analogue (especially in light of the polynomial upper bounds that we derived in Section 3), is the following.

Corollary 4.17. *For all $n \geq 2$, there exist n -vertex graphs X and Y for which there exists a connected component of $\text{FS}(X, Y)$ such that the friends-and-strangers Markov chain of X and Y , when restricted to this component, has mixing time (in total variation distance) which is $e^{\Omega(n)}$.*

In other words, for this variant of the interchange process in which we may further forbid certain pairs of particles from swapping places with each other, it is possible to fix restrictions between particles in such a way that the mixing time of the underlying Markov chain is exponential in the size of the graph on which the process occurs. This is in stark contrast to the aforementioned polynomial upper bounds regarding rapidly mixing Markov chains.

5. Open Questions and Future Directions

Theorem 1.3 of this paper proves that diameters of connected components of friends-and-strangers graphs may grow exponentially in the size of their input graphs. There are many other interesting questions concerning distance and diameter that remain unresolved by this article.

5.1. Other Choices of Fixed Graphs

In Section 3, we fixed X to be from a particular class of graphs, and derived bounds on the maximal diameter of a connected component $\text{FS}(X, Y)$. Of course, we could pursue similar

¹³The interchange process is usually posed as a continuous-time stochastic process by assigning independent point processes on the positive half-line to the edges of X and transposing the particles upon the vertices incident to a given edge at the points of its corresponding process. One can certainly adapt Definition 4.16 to accommodate for such differences in the presentation of the model.

inquiries for other choices of X . One natural choice would be to take $X = \text{Star}_n$. It is known (see [BJL⁺23]) that the diameter of any component of $\text{FS}(\text{Star}_n, K_n)$ is at most $\frac{3}{2}n + O(1)$, but to our knowledge, there are no known bounds on the maximum diameter for general Y . We also remark that it may be possible to extract a bound on the maximum diameter of a component of $\text{FS}(\text{Star}_n, Y)$ for biconnected graphs Y by tracing the arguments in [Wil74].

5.2. Improvements

For much of our discussion in Subsection 3.3, we were primarily interested in showing that the maximum diameter of a connected component of $\text{FS}(\text{Cycle}_n, Y)$ was polynomially bounded (in the sense of Question 1.2), rather than achieving tight asymptotic statements. It would be desirable to improve these results, toward which we pose the following conjectures. We mention that generalizing the lattice-theoretic methods of [Pro21], which are of a very different flavor than the arguments presented here, might lead to the resolution of Conjecture 5.2. We note that if Conjecture 5.1 were settled, then tracing the proof of Corollary 3.10 would immediately settle Conjecture 5.2. On the other hand, if Conjecture 5.2 were settled, then an adaptation of the proof of Theorem 3.7 would sharpen Theorem 3.9 and lead to an $O(n^3)$ bound on the diameter of any connected component of $\text{FS}(\text{Cycle}_n, Y)$.

Conjecture 5.1. The maximum diameter of a connected component of $\text{FS}(\text{Cycle}_n, Y)$ is $O(n^2)$.

Conjecture 5.2. For an n -vertex graph G and two acyclic orientations $\alpha, \alpha'' \in \text{Acyc}(G)$ that are double-flip equivalent, it is possible to go from α to α'' in $O(n^2)$ double-flips.

In another direction, Theorem 1.3 states that for all $n \geq 2$, there exist n -vertex X and Y such that the maximum diameter of a connected component of $\text{FS}(X, Y)$ is $e^{\Omega(n)}$. It is unclear how close this is to the truth. As a first step, we pose the following problem.

Question 5.3. For n -vertex graphs X and Y , does there exist a nontrivial upper bound (in terms of n) on the maximum diameter of a component of $\text{FS}(X, Y)$?

We briefly clarify what we mean by a nontrivial upper bound in Question 5.3. Let $\mathcal{D}(n)$ denote the maximum possible diameter of a connected component of $\text{FS}(X, Y)$ when X and Y are n -vertex graphs. By Theorem 1.3 for the lower bound and Stirling's approximation applied to $n!$ for the upper bound, we observe that

$$e^{\Omega(n)} = \mathcal{D}(n) \leq e^{(1-o(1))n \log n}.$$

Thus, our understanding of $\mathcal{D}(n)$ is tight up to a logarithmic factor in the exponent. Any improvement over this naive upper bound, or confirmation that this upper bound is essentially the truth, would be highly desirable. In particular, we ask the following more precise question. Indeed, given the preceding discussion, Question 5.4 is the natural next target.

Question 5.4. Is it true that $\mathcal{D}(n) = e^{O(n)}$?

We propose one final problem in this subsection which we would especially like to see resolved.

Problem 5.5. Find a shorter (perhaps via non-constructive¹⁴ means) proof of Theorem 1.3.

5.3. Connected Friends-and-Strangers Graphs

The proof of Theorem 1.3 relied heavily on characterizing all vertices of $\text{FS}(X_L, Y_L)$ in the same connected component of σ_s . It is thus natural to ask Question 1.2 in the setting where $\text{FS}(X, Y)$ is assumed to be connected, which was separately raised by Defant and Kravitz.

Question 5.6 ([DK21, Subsection 7.3]). Does there exist an absolute constant $C > 0$ such that for all n -vertex graphs X and Y with $\text{FS}(X, Y)$ connected, it holds that $\text{diam}(\text{FS}(X, Y))$ is $O(n^C)$?

If Question 5.6 holds in the negative, then settling it will likely require very different techniques and paradigms than those which were developed in this article. Indeed, the proof of the negative result for Question 1.2 relies heavily on “rigging” the configurations that lie in a particular connected component of $\text{FS}(X_L, Y_L)$, which allows us to argue that two particular configurations (namely, σ_s and σ_f) are necessarily very far apart. Such a strategy is not applicable if we require $\text{FS}(X, Y)$ to be connected. Additionally, by Proposition 2.3(3), we can assume (without loss of generality) that X is biconnected under this setting. Theorem 3.7 already gives a positive result for Cycle_n , the “simplest” biconnected graph (e.g., the n -vertex cycle has the smallest Betti number amongst all n -vertex biconnected graphs: see [Whi31, Theorem 19], which might lend itself to an inductive argument) and for K_n , the most “complicated” (K_n has the largest Betti number amongst all n -vertex biconnected graphs). Furthermore, the constructions X_L and Y_L contain cut vertices which hold central roles in the proofs of the intermediate propositions (namely, vertices on the paths $\mathcal{P}_a^\ell, \mathcal{P}_b^\ell$ for X_L , and the knob vertices $\kappa_a^\ell, \kappa_b^\ell$ in Y_L).

In another direction, a negative answer to Question 5.6 implies the existence of long paths in the connected graph $\text{FS}(X, Y)$. The following result shows that the extreme end of this is not possible.

Proposition 5.7. For $n \geq 4$, $\text{FS}(X, Y)$ is not isomorphic to a tree on $n!$ vertices (e.g., $\text{Path}_{n!}$) or a tree on $n!$ vertices with one edge appended (e.g., $\text{Cycle}_{n!}$).

Proof. The number of edges of $\text{FS}(X, Y)$ is $|E(X)| \cdot |E(Y)| \cdot (n - 2)!$, while this is $n! - 1$ and $n!$ for a tree on $n!$ vertices and a tree with one edge appended on $n!$ vertices, respectively. Notice that $|E(X)| \cdot |E(Y)| \cdot (n - 2)!$ is divisible by 2 while $n! - 1$ is not, so $\text{FS}(X, Y)$ cannot be isomorphic to a tree on $n!$ vertices. Assume $\text{FS}(X, Y)$ is isomorphic to a tree with an edge appended to it, so $|E(X)| \cdot |E(Y)| \cdot (n - 2)! = n!$, or $|E(X)| \cdot |E(Y)| = n(n - 1)$. Then X and Y must both be connected, so that (without loss of generality) $|E(X)| = n$ and $|E(Y)| = n - 1$, so Y is a tree. Due to Proposition 2.3(3), X is biconnected, so necessarily $X = \text{Cycle}_n$. But $|E(\bar{Y})| = \binom{n}{2} - (n - 1)$, contradicting Theorem 2.6, which gives $|E(\bar{Y})| \leq n - 1$. \square

¹⁴Certainly, a proof of Theorem 1.3 using non-constructive techniques would be a novel contribution. In another direction, recall that the central idea behind Section 4 was to construct an exponentially increasing recursive sequence of swaps in such a way that executing this sequence of swaps is necessary in order to reach one configuration from another. A constructive proof which either proceeds via a similar paradigm with a construction that is more amenable to analysis or leverages different ideas altogether would also be of interest.

5.4. Probabilistic Problems

In a different direction, we may study notions of distance in friends-and-strangers graphs when we take X and Y to be random graphs. We propose the following problem; we leave the meaning of “small diameter” up to interpretation.

Problem 5.8. Let X and Y be independently-chosen random graphs from $\mathcal{G}(n, p)$. Find conditions on p (in terms of n) which guarantee that every connected component of $\text{FS}(X, Y)$ has small diameter with high probability.

We also restate a problem of this kind proposed by [ADK23].

Problem 5.9 ([ADK23, Problem 7.9]). Obtain estimates (in terms of n and p) for the expectation of the maximum diameter of a connected component in $\text{FS}(X, Y)$ when X and Y are independently-chosen random graphs from $\mathcal{G}(n, p)$.

In a manner analogous to how we fixed one of the two graphs X and Y in Section 3 and studied the resulting variant of Question 1.2 before addressing the more global question, it may be insightful to first fix (without loss of generality) X to be a particular kind of graph and study the variants of Problems 5.8 and 5.9 which only take Y to be drawn from $\mathcal{G}(n, p)$. The graphs we studied in Section 3 (complete graphs, paths, and cycles) may also serve as natural starting points here.

5.5. Complexity

As the literature on the token swapping problem suggests, computing exact distances between two configurations in $\text{FS}(X, Y)$ and the maximum diameter of a component of $\text{FS}(X, Y)$, under mild assumptions on X and Y , seems to be intractable. We might thus study distances and diameters in friends-and-strangers graphs from the perspective of complexity theory. We discuss one possible direction of study along these lines here. We start by introducing a decision problem which encapsulates finding the shortest swap sequence between two configurations.

Definition 5.10. In an instance of the *distance problem*, we are given graphs X and Y on n vertices, configurations $\sigma, \tau \in V(\text{FS}(X, Y))$, and a positive integer K , and want to know if $d(\sigma, \tau) \leq K$.

This problem has been studied in many restricted contexts. If we proceed under the assumption that $Y = K_n$, the distance problem is known to be PSPACE-complete [Jer85], APX-hard [MNO⁺16], and $W[1]$ -hard when parametrized by the shortest number of swaps [BMR18]. Furthermore, it is NP-hard when we impose certain additional restrictions, such as when we take X to be a tree and $Y = K_n$ [ADK⁺22]. It might be fruitful to study the complexity of this problem at the level of generality proposed by Definition 5.10.

A natural follow-up to Definition 5.10 is to ask for worst-case distances between two configurations, which corresponds to the maximum diameter of a component of $\text{FS}(X, Y)$.

Definition 5.11. In an instance of the *diameter problem*, we are given graphs X and Y on n vertices and a positive integer K , and want to know if the maximum diameter of a component of $\text{FS}(X, Y)$ is at most K .

The literature suggests that the diameter problem has not been studied as thoroughly as the distance problem, even when assuming that one of the two graphs is complete. Towards bridging this gap, we pose the following primitive question and problem.

Question 5.12. Is the diameter problem in EXPSPACE? If so, is it EXPSPACE-complete? What changes if we fix $Y = K_n$?

Problem 5.13. Find assumptions on X , Y , and K which guarantee that the diameter problem (under these assumptions) is in PSPACE.

Notably, even simpler decision problems than those proposed in Definitions 5.10 and 5.11 seem to be poorly understood. For instance, we may consider the following decision problems.

Definition 5.14. In an instance of the *component problem*, we are given n -vertex graphs X and Y and configurations $\sigma, \tau \in V(\text{FS}(X, Y))$, and want to know if σ and τ lie in the same connected component of $\text{FS}(X, Y)$.

Definition 5.15. In an instance of the *connectivity problem*, we are given n -vertex graphs X and Y , and want to know if $\text{FS}(X, Y)$ is connected.

We may think of the component problem and the connectivity problem as, respectively, the simplest instances of the distance problem (asking whether $d(\sigma, \tau)$ is finite) and of the diameter problem (asking whether the diameter of $\text{FS}(X, Y)$, when we do not restrict to connected components, is finite). To our knowledge, an understanding of the complexity of the component problem and the connectivity problem remains open, though the results in [ABC⁺23] address these problems when studying friends-and-strangers graphs with multiplicities, as elaborated in [Mil24].

In a different direction, [AK89, VP95, YDI⁺15] independently found 2-approximation algorithms for determining the distance between two configurations in $\text{FS}(X, K_n)$ when X is a tree. Recall from the proof of Proposition 3.2 that for any $\sigma, \tau \in \text{FS}(\text{Path}_n, Y)$ in the same connected component, $d(\sigma, \tau) = \text{inv}(\sigma, \tau)$, and an algorithm which exactly computes the distance between any two configurations in $\text{FS}(\text{Path}_n, Y)$ is one which, starting from σ , reverses a (σ, τ) -inversion at every step. These two observations naturally suggest the following problem, which one can also pursue by replacing Cycle_n with a different fixed graph.

Problem 5.16. Find, under the most general assumptions on Y possible, an $O(1)$ -approximation algorithm for computing the distance between two configurations in $\text{FS}(\text{Cycle}_n, Y)$.

Acknowledgements

I would like to thank Professor Joseph Gallian for organizing the Duluth REU and for giving me a chance to participate in his 2021 program, and I am deeply grateful to Colin Defant and Noah Kravitz for suggesting these problems to me and for many helpful conversations over the course of that summer. In particular, I thank Colin Defant for carefully reading through a draft of this article, providing many productive comments on the manuscript. I also thank Yelena

Mandelstam for reviewing a draft of this work and providing constructive feedback. Finally, I would like to sincerely thank the two anonymous referees for thoroughly reviewing an earlier draft of this work, catching many typos and grammatical mistakes, and providing many detailed comments and suggestions which have vastly improved this article’s clarity. Most notably, I am deeply indebted to the anonymous referee who proposed a simple way to sharpen the main result of this paper and who suggested that I consider implications of Theorem 1.3 to the study of the interchange process, which very quickly led to substantially improved results (both qualitatively and quantitatively) in this article.

References

- [ABC⁺23] Robert M Alaniz, Josh Brunner, Michael Coulombe, Erik D Demaine, Jenny Diomidova, Timothy Gomez, Elise Grizzell, Ryan Knobel, Jayson Lynch, Andrew Rodriguez, et al. Complexity of reconfiguration in surface chemical reaction networks. In *29th International Conference on DNA Computing and Molecular Programming (DNA 29)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2023.
- [AD86] David Aldous and Persi Diaconis. Shuffling cards and stopping times. *The American Mathematical Monthly*, 93(5):333–348, 1986. doi:10.2307/2323590.
- [ADK⁺22] Oswin Aichholzer, Erik D Demaine, Matias Korman, Anna Lubiw, Jayson Lynch, Zuzana Masárová, Mikhail Rudoy, Virginia Vassilevska Williams, and Nicole Wein. Hardness of token swapping on trees. In *30th Annual European Symposium on Algorithms: ESA 2022*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.
- [ADK23] Noga Alon, Colin Defant, and Noah Kravitz. Typical and extremal aspects of friends-and-strangers graphs. *Journal of Combinatorial Theory, Series B*, 158(1):3–42, 2023. doi:10.1016/j.jctb.2022.03.001.
- [AF02] David Aldous and Jim Fill. Reversible markov chains and random walks on graphs. 2002. URL: <https://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- [AK89] Sheldon B. Akers and Balakrishnan Krishnamurthy. A group-theoretic model for symmetric interconnection networks. *IEEE Transactions on Computers*, 38(4):555–566, 1989. doi:10.1109/12.21148.
- [AK13] Gil Alon and Gady Kozma. The probability of long cycles in interchange processes. *Duke Mathematical Journal*, 162(9):1567–1585, 2013. doi:10.1215/00127094-2266018.
- [Ald83] David Aldous. Random walks on finite groups and rapidly mixing markov chains. In *Séminaire de Probabilités XVII 1981/82: Proceedings*, pages 243–297. Springer, 1983. doi:10.1007/bfb0068322.
- [Alo21] Noga Alon. Random friends walking on random graphs, 2021. URL: <https://www.youtube.com/watch?v=uheP2ctJx10>.
- [Ang03] Omer Angel. Random infinite permutations and the cyclic time random walk. In *DMTCS Proceedings vol. AC, Discrete Random Walks (DRW’03)*, pages 9–16.

- Discrete Mathematics and Theoretical Computer Science, 2003. doi:10.46298/dmtcs.3342.
- [Ban22] Kiril Bangachev. On the asymmetric generalizations of two extremal questions on friends-and-strangers graphs. *European Journal of Combinatorics*, 104, 2022. doi:10.1016/j.ejc.2022.103529.
- [BD06] Nathanaël Berestycki and Rick Durrett. A phase transition in the random transposition random walk. *Probability Theory and Related Fields*, 136:203–233, 2006. doi:10.1007/s00440-005-0479-7.
- [BE78] Béla Bollobás and Stephen E Eldridge. Packings of graphs and applications to computational complexity. *Journal of Combinatorial Theory, Series B*, 25(2):105–124, 1978. doi:10.1016/0095-8956(78)90030-8.
- [BGYW19] Tatiana Brailovskaya, Gokul Gowri, Sean Yu, and Erik Winfree. Reversible computation using swap reactions on a surface. In *DNA Computing and Molecular Programming: 25th International Conference, DNA 25, Seattle, WA, USA, August 5–9, 2019, Proceedings 25*, pages 174–196. Springer, 2019.
- [BJL⁺23] Ahmad Biniiaz, Kshitij Jain, Anna Lubiw, Zuzana Masárová, Tillmann Miltzow, Debajyoti Mondal, Anurag Murty Naredla, Josef Tkadlec, and Alexi Turcotte. Token swapping on trees. *Discrete Mathematics & Theoretical Computer Science*, 24(2), 2023. doi:10.46298/dmtcs.8383.
- [BJS17] Béla Bollobás, Svante Janson, and Alex Scott. Packing random graphs and hypergraphs. *Random Structures & Algorithms*, 51(1):3–13, 2017. doi:10.1002/rsa.20673.
- [BK23] Florestan Brunck and Matthew Kwan. Books, hallways and social butterflies: a note on sliding block puzzles. 2023. arXiv:2303.09459.
- [BMR18] Édouard Bonnet, Tillmann Miltzow, and Paweł Rzkażewski. Complexity of token swapping and its variants. *Algorithmica*, 80:2656–2682, 2018. doi:10.1007/s00453-017-0387-0.
- [BR99] Christopher L Barrett and Christian M Reidys. Elements of a theory of computer simulation i: sequential ca over random graphs. *Applied Mathematics and Computation*, 98(2-3):241–259, 1999. doi:10.1016/s0096-3003(97)10166-7.
- [Cay49] Arthur Cayley. Lxxvii. note on the theory of permutations. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 34(232):527–529, 1849. doi:10.1080/14786444908646287.
- [Che10] Beifang Chen. Orientations, lattice polytopes, and group arrangements i: chromatic and tension polynomials of graphs. *Annals of Combinatorics*, 13:425–452, 2010. doi:10.1007/s00026-009-0037-6.
- [CLR10] Pietro Caputo, Thomas Liggett, and Thomas Richthammer. Proof of aldous’ spectral gap conjecture. *Journal of the American Mathematical Society*, 23(3):831–851, 2010. doi:10.1090/s0894-0347-10-00659-4.

- [DDLW24] Colin Defant, David Dong, Alan Lee, and Michelle Wei. Connectedness and cycle spaces of friends-and-strangers graphs. *Discussiones Mathematicae Graph Theory*, 44(3), 2024. doi:10.7151/dmgt.2492.
- [DK21] Colin Defant and Noah Kravitz. Friends and strangers walking on graphs. *Combinatorial Theory*, 1(0), 2021. doi:10.5070/c61055363.
- [DMR16] Mike Develin, Matthew Macauley, and Victor Reiner. Toric partial orders. *Transactions of the American Mathematical Society*, 368(4):2263–2287, 2016. doi:10.1090/tran/6356.
- [DR18] Erik D Demaine and Mikhail Rudoy. A simple proof that the $(n^2 - 1)$ -puzzle is hard. *Theoretical Computer Science*, 732:80–84, 2018. doi:10.1016/j.tcs.2018.04.031.
- [DS81] Persi Diaconis and Mehrdad Shahshahani. Generating a random permutation with random transpositions. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 57(2):159–179, 1981. doi:10.1007/bf00535487.
- [DSC93] Persi Diaconis and Laurent Saloff-Coste. Comparison techniques for random walk on finite groups. *The Annals of Probability*, 21(4):2131–2156, 1993. doi:10.1214/aop/1176989013.
- [ES23] Dor Elboim and Allan Sly. Infinite cycles in the interchange process in five dimensions. 2023. arXiv:2211.17023.
- [Ham15] Alan Hammond. Sharp phase transition in the random stirring model on trees. *Probability Theory and Related Fields*, 161:429–448, 2015. doi:10.1007/s00440-013-0543-7.
- [HS21] Jonathan Hermon and Justin Salez. The interchange process on high-dimensional products. *Annals of Applied Probability*, 31(1), 2021. doi:10.1214/20-aap1583.
- [Jeo22] Ryan Jeong. On structural aspects of friends-and-strangers graphs. To appear in the *Annals of Combinatorics*, 2022. arXiv:2203.10337.
- [Jeo23] Ryan Jeong. Bipartite friends and strangers walking on bipartite graphs. 2023. arXiv:2309.03848.
- [Jer85] Mark R Jerrum. The complexity of finding minimum-length generator sequences. *Theoretical Computer Science*, 36:265–289, 1985. doi:10.1016/0304-3975(85)90047-7.
- [Jon12] Johan Jonasson. Mixing times for the interchange process. *ALEA*, 9(2):667–683, 2012.
- [Kim16] Dohan Kim. Sorting on graphs by adjacent swaps using permutation groups. *Computer Science Review*, 22:89–105, 2016. doi:10.1016/j.cosrev.2016.09.003.
- [KMS84] D. Kornhauser, G. Miller, and P. Spirakis. Coordinating pebble motion on graphs, the diameter of permutation groups, and applications. In *25th Annual Symposium on Foundations of Computer Science, 1984*, pages 241–250, 1984.

- [KO09] Daniela Kühn and Deryk Osthus. The minimum degree threshold for perfect graph packings. *Combinatorica*, 29(1):65–107, 2009. doi:10.1007/s00493-009-2254-3.
- [KSY19] Jun Kawahara, Toshiki Saitoh, and Ryo Yoshinaka. The time complexity of permutation routing via matching, token swapping and a variant. *Journal of Graph Algorithms and Applications*, 23(1):29–70, 2019. doi:10.7155/jgaa.00483.
- [Lee22] Alan Lee. Connectedness in friends-and-strangers graphs of spiders and complements. 2022. arXiv:2210.04768.
- [Lov93] L Lovász. Random walks on graphs: a survey. *J. Bolyai Math. Soc.*, 2:1–46, 1993.
- [Mat88] Peter Matthews. A strong uniform time for random transpositions. *Journal of Theoretical Probability*, 1:411–423, 1988. doi:10.1007/bf01048728.
- [Mil24] Aleksa Milojević. Connectivity of old and new models of friends-and-strangers graphs. *Advances in Applied Mathematics*, 155:102668, 2024. doi:10.1016/j.aam.2023.102668.
- [MM11] Matthew Macauley and Henning S Mortveit. Posets from admissible coxeter sequences. *The Electronic Journal of Combinatorics*, 18(1):P197, 2011. doi:10.37236/684.
- [MNO⁺16] Tillmann Miltzow, Lothar Narins, Yoshio Okamoto, Günter Rote, Antonis Thomas, and Takeaki Uno. Approximation and hardness of token swapping. In *24th Annual European Symposium on Algorithms*, volume 57, pages 66:1–66:15, 2016.
- [Naa00] Michael Naatz. The graph of linear extensions revisited. *SIAM Journal on Discrete Mathematics*, 13(3):354–369, 2000. doi:10.1137/s0895480199352609.
- [Par15] Ian Parberry. Solving the $(n^2 - 1)$ -puzzle with $\frac{8}{3}n$ expected moves. *Algorithms*, 8(3):459–465, 2015. doi:10.3390/a8030459.
- [Pre86] Oliver Pretzel. On reorienting graphs by pushing down maximal vertices. *Order*, 3(2):135–153, 1986. doi:10.1007/bf00390104.
- [Pro21] James Propp. Lattice structure for orientations of graphs. 2021. arXiv:math/0209005.
- [PV90] Frederick J Portier and Theresa P Vaughan. Whitney numbers of the second kind for the star poset. *European Journal of Combinatorics*, 11(3):277–288, 1990. doi:10.1016/s0195-6698(13)80127-8.
- [Rei98] CM Reidys. Acyclic orientations of random graphs. *Advances in Applied Mathematics*, 21(2):181–192, 1998. doi:10.1006/aama.1998.0595.
- [Sch05] Oded Schramm. Compositions of random transpositions. *Israel Journal of Mathematics*, 147:221–243, 2005. doi:10.1007/bf02785366.
- [Spe09] David Speyer. Powers of coxeter elements in infinite groups are reduced. *Proceedings of the American Mathematical Society*, 137(4):1295–1302, 2009. doi:10.1090/s0002-9939-08-09638-x.

- [SS78] Norbert Sauer and Joel Spencer. Edge disjoint placement of graphs. *Journal of Combinatorial Theory, Series B*, 25(3):295–302, 1978. doi:10.1016/0095-8956(78)90005-9.
- [Sta12] Richard P Stanley. An equivalence relation on the symmetric group and multiplicity-free flag h-vectors. *Journal of Combinatorics*, 3(3):277–298, 2012. doi:10.4310/joc.2012.v3.n3.a2.
- [vBSY16] Anke van Zuylen, James Bieron, Frans Schalekamp, and Gexin Yu. A tight upper bound on the number of cyclically adjacent transpositions to sort a permutation. *Information Processing Letters*, 116(11):718–722, 2016. doi:10.1016/j.ipl.2016.06.006.
- [VP95] Theresa P Vaughan and Frederick J Portier. An algorithm for the factorization of permutations on a tree. *Journal of Combinatorial Mathematics and Combinatorial Computing*, 18:11–31, 1995.
- [WC23] Lanchao Wang and Yaojun Chen. Connectivity of friends-and-strangers graphs on random pairs. *Discrete Mathematics*, 346(3):113266, 2023. doi:10.1016/j.disc.2022.113266.
- [Whi31] Hassler Whitney. Non-separable and planar graphs. *Proceedings of the National Academy of Sciences*, 17(2):125–127, 1931. doi:10.1073/pnas.17.2.125.
- [Wil74] Richard M Wilson. Graph puzzles, homotopy, and the alternating group. *Journal of Combinatorial Theory, Series B*, 16(1):86–96, 1974. doi:10.1016/0095-8956(74)90098-7.
- [Wil04] David Bruce Wilson. Mixing times of lozenge tiling and card shuffling markov chains. *The Annals of Applied Probability*, 14(1):274–325, 2004. doi:10.1214/aop/1075828054.
- [WLC23] Lanchao Wang, Junying Lu, and Yaojun Chen. Connectedness of friends-and-strangers graphs of complete bipartite graphs and others. *Discrete Mathematics*, 346(8):113499, 2023. doi:10.1016/j.disc.2023.113499.
- [Yan11] Chao Yang. Sliding puzzles and rotating puzzles on graphs. *Discrete Mathematics*, 311(14):1290–1294, 2011. doi:10.1016/j.disc.2011.03.011.
- [Yap88] HP Yap. Packing of graphs-a survey. *Discrete Mathematics*, 72(1-3):395–404, 1988. doi:10.1016/0012-365x(88)90232-4.
- [YDI⁺15] Katsuhisa Yamanaka, Erik D Demaine, Takehiro Ito, Jun Kawahara, Masashi Kiyomi, Yoshio Okamoto, Toshiki Saitoh, Akira Suzuki, Kei Uchizawa, and Takeaki Uno. Swapping labeled tokens on graphs. *Theoretical Computer Science*, 586:81–94, 2015. doi:10.1016/j.tcs.2015.01.052.
- [Yus07] Raphael Yuster. Combinatorial and computational aspects of graph packing and graph decomposition. *Computer Science Review*, 1(1):12–26, 2007. doi:10.1016/j.cosrev.2007.07.002.
- [Zhu24] Honglin Zhu. The connectivity of friends-and-strangers graphs on complete multipartite graphs. 2024. arXiv:2307.08121.