Detecting Graphs with Induced Cycles in Limited-Bandwidth Communication Networks

by

Timothy Zapp

A thesis submitted to
The Faculty of Graduate Studies of
The University of Manitoba
in partial fulfillment of the requirements
of the degree of

Master of Science

Department of Computer Science
The University of Manitoba
Winnipeg, Manitoba, Canada
August 2023

Thesis advisor

Avery Miller

Author

Timothy Zapp

Detecting Graphs with Induced Cycles in Limited-Bandwidth Communication Networks

Abstract

The subgraph detection problem is a fundamental problem in distributed algorithms that asks the nodes of a communication network G to determine if that communication network contains a specific graph P as a subgraph. Recent work in this area has largely been focused on the case where the nodes have limited bandwidth and thus cannot communicate arbitrarily, which is part of a larger movement in distributed algorithms to reduce the amount of communication that our algorithms require.

We show that, if P is a graph containing an induced cycle of length at least 4, then detecting a rainbow copy of P (in terms of vertex colourings or edge colourings) takes at least $\Omega(n/\log n)$ in the CONGEST model. Motivated by this result, we go on to show that, if P is a necklace graph containing an odd cycle of length at least 5, then detecting a copy of P takes at least $\Omega(n/\log n)$ rounds in the CONGEST model. Finally, we discuss techniques for proving similar results when P is a necklace graph containing all even cycles or a 3-cycle, and highlight the difficulties in doing so.

Contents

Acknowledgments

I would like to begin by thanking Avery Miller and my committee for all of their hard work supervising the writing of this thesis. This would not have been possible without them. I would also like to thank my family and friends, who have been a source of endless support. Last, but certainly not least, I would like to thank Pumpkin, for the unwavering companionship.

This thesis is dedicated to those that perished on the climb.

Chapter 1

Introduction

The study of congested models of distributed computing is motivated by the desire to reduce the bandwidth costs of our distributed systems. In particular, we have recently seen the size and complexity of our distributed systems increase drastically, due to an increase in users and in network-connected devices, which causes the communication cost of algorithms on these systems to increase as well. Additionally, individual network usage has increased due to factors such as an increase in streaming services (tv, music, etc) and an increase in remote ("cloud") storage of data. However, there are physical and technological limitations on moving large amounts of data in a short amount of time. This means that it is no longer realistic to study theoretical models of networks that allow arbitrarily large message sizes to be sent in one round, as was allowed in earlier work in the literature. We must consider the effect that restricting bandwidth has on our algorithmic techniques to ensure that our networks remain responsive. This has resulted in a re-examination of fundamental distributed problems, and subgraph detection is a fundamental problem in distributed algorithms and graph algorithms. In the distributed case, this is motivated by the fact that we often want to monitor properties of the network topology, especially as the topology may change over time. Detecting cliques, for example, can be used to find highly connected portions of the network, allowing quick information dissemination. To handle changes over time, the presence of cycles, for instance, can be used to ensure that we have redundancy in the network, so that a node leaving the network does not disconnect the nodes in the cycle. In this thesis we prove lower bounds on the cost of algorithms for the subgraph detection problem as these results help provide insight into whether our algorithms are as efficient as possible or if there is still room for improvement and, in the latter case, where we might look to improve existing algorithms.

1.1 Preliminaries

Here we introduce the mathematical concepts and notation that will be used throughout this thesis. First, for each positive integer n, we define the set $[n] = \{1, 2, \ldots, n-1, n\}$. For any binary string X, we refer to the i-th bit of X using the notation X_i. Let S, T be sets and let $f : S \to T$ be a function. For any $T' \subseteq T$, we define the set $f^{-1}(T') = \{s \in S : f(s) \in T'\}$.

We model networks using graphs. A *graph* is an ordered pair $G = (V, E)$, where V is a set and E is a set of 2-element subsets of V. The elements of V are called *vertices* (sometimes *nodes*) and the elements of E are called *edges*. We often visualize a graph by drawing the vertices as points and the edges as curves so that, if $\{u, v\} \in E$, then there is a curve connecting the point representing u and the point representing v. For convenience, we will write an edge $\{u, v\}$ as uv. Finally, for the sake of clarity, when we are referring to a graph G we will often refer to its vertex set as $V(G)$ and its edge set as $E(G)$.

If G is a graph and $v \in V(G)$, the *neighbourhood* of v in G is the set $N_G(v) = \{u \in V(G) : uv \in E(G)\}$. If $uv \in E(G)$, we say that u and v are *adjacent* and that u and v are *incident* to the edge uv.

A graph H is a *subgraph* of a graph G if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. Let $S \subseteq V(G)$. Then the *subgraph of* G *induced by* S, denoted by $G[S]$, is the graph with vertex set S and edges so that $uv \in E(G[S])$ if and only if $uv \in E(G)$.

Let G and H be graphs. G is said to *contain a copy of* H or *contain* H *as a subgraph* if there is an injective function $\phi : V(H) \to V(G)$ so that, for each $uv \in E(H)$, $\phi(u)\phi(v) \in E(G)$. We call such a mapping ϕ the *subgraph mapping*. Two graphs G and H are *isomorphic* if G contains H as a subgraph, the subgraph mapping ϕ is surjective, and $uv \in E(H)$ if and only if $\phi(u)\phi(v) \in E(G)$. In this case, we call the function ϕ an *isomorphism*. Intuitively, if two graphs are isomorphic, then they are structurally the same graph but with potentially different labels for their vertices.

Next, we describe some classes of graphs that are important for the statements and proofs of our results. A *path* of length n is a graph P_{n+1} with vertices $v_0, v_1, \ldots, v_n$ and edges $v_{i-1}v_i$, for each $i \in [n]$. Similarly, a *cycle* of length n is a graph C_n with vertices $v_0, v_1, \ldots, v_{n-1}$ and edges $v_{(i-1) \bmod n} v_{i \bmod n}$, for each $i \in [n]$. A *complete graph* on n vertices is a graph K_n with n vertices such that every two vertices are adjacent. A graph G is called *bipartite* if there is a partition of its vertex set into parts A, B such that $V(G) = A \cup B$ and each edge of G is between a vertex in A and a vertex in B. For

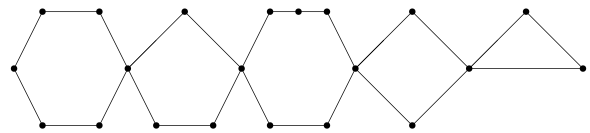

Figure 1.1: An example of a necklace graph with 5 cycles.

any positive integers n, m, the *complete bipartite* graph $K_{n,m}$ is the graph with vertex set $\{a_1, a_2, \ldots, a_n, b_1, b_2, \ldots, b_m\}$ and edge set $\{a_i b_j : i \in [n], j \in [m]\}$. Intuitively, a complete bipartite graph is a bipartite graph that contains every possible edge.

Let G be a graph that contains P_n as a subgraph with subgraph mapping ϕ. Then we can write $\phi(V(P_n))$ as a finite sequence of vertices of G so that each pair of consecutive vertices in the sequence are adjacent. Then the first and last vertices in this sequence are called the *endpoints* of the path. The vertices of $\phi(V(P_n))$ that are not endpoints are called *internal nodes* of the path. If u and v are the endpoints of a path, we often say that the path is *from u to v*. A graph G is *connected* if, for any pair of vertices $u, v \in V(G)$, there is a copy of a path in G (often just called a *path in G*, for conciseness) from u to v. If a graph G contains two paths, the copies are said to be *disjoint* if they share no vertices and *internally disjoint* if they share no internal nodes (but potentially share endpoints).

One focus of our work will involve the class of necklace graphs. A *necklace graph with m cycles* is a connected graph so that exactly $m-1$ vertices belong to exactly two cycles and each of the remaining vertices belong to exactly one cycle. Intuitively, we can imagine a necklace graph as a sequence of cycles $C_1, \ldots, C_m$ that have been "glued" together to form a path, and each pair of neighbouring cycles has exactly one vertex in common. See Figure 1.1 for an example of a necklace graph.

Finally, we will need some concepts from extremal graph theory. Let H be a graph and n be a positive integer. Then, we define the set $\mathcal{G}_{n,H}$ to be the set of all graphs with n vertices that do not contain H as a subgraph. Then the *Túran number* (also called the *extremal number*) for the pair n, H is $\text{ex}(n, H) = \max\{|E(G)| : G \in \mathcal{G}_{n,H}\}$. In other words, the Túran number for n, H is the maximum number of edges in a graph with n vertices that does not contain H as a subgraph. Each graph $G \in \mathcal{G}_{n,H}$ with exactly $\text{ex}(n, H)$ edges is called an *extremal graph* for the pair n, H. We will sometimes also call a graph G an extremal graph for n, H if it has n vertices, $\Theta(\text{ex}(n, H))$ edges, and does not contain H as a subgraph, as asymptotic bounds are often good enough for our purposes.

1.2 Computational Model

Models of congested networks have recently been the subject of significant research interest in theoretical distributed computing. Each of the congestion models assumes that we have a network of n processes (often referred to as nodes). Each node has a unique ID, which is assumed to be encoded using a $O(\log n)$-bit binary string. The *communication graph* C is a graph on n vertices so that the vertices of C represent the nodes and the edges of C represent the communication links in the network. We assume that time proceeds in discrete, synchronous rounds. Each round consists of three phases. First, each node may perform unlimited local computation. Then, each node sends at most one (potentially different) message to each of its neighbours in C. Finally, each node receives any messages that it was sent by its neighbours in this round. We define a parameter b, called the *bandwidth*, so that each message sent can consist of at most b bits. In congested network models, it is usually assumed that

$b = \log n$ or $b \in O(\log n)$ so that, at minimum, each node can send its own ID in each message. All results discussed in this thesis hold under either assumption.

When working with these models we are often concerned with graph problems, so we define the *input graph* G to be a graph on n vertices so that each node receives a distinct vertex $v \in V(G)$ along with its neighbourhood $N_G(v)$ as input. In other words, each node is identified with a vertex v and knows v and all of v's neighbours. The two most prominent models in this family are the CONGEST model, in which we assume that $G = C$, and the CLIQUE model, also called the *congested clique* model, in which we assume that $C = K_n$. Intuitively it seems that the CLIQUE model is much more powerful than the CONGEST model, as there are (usually) more communication links in C and so more information can be exchanged throughout the network in each round. This is backed up somewhat by the fact that very few non-trivial lower bounds are known for the CLIQUE model. Drucker et al. [11] back this up more formally by showing that establishing lower bounds for decision problems in the CLIQUE model would solve open problems in circuit complexity, partially explaining why showing lower bounds for this model is difficult. On the other hand, any lower bound that holds in the CLIQUE model trivially also holds in the CONGEST model, since any algorithm for the CONGEST model can be run without modification in the CLIQUE model by ignoring any edges that are present in C but not in G.

The primary metric that we will use to measure the performance of algorithms in these models is the round complexity. Let $\mathcal{A}$ be an algorithm for solving some graph problem in this model. Now, each node must terminate their execution after a finite number of rounds and specify an output value. Let C be a communication graph on n nodes, let $v \in V(C)$, and let G be an input graph (which means it has n vertices). Then let $r(v, \mathcal{A}, C, G)$ be the number of rounds after which v terminates when running the algorithm $\mathcal{A}$ with communication graph C and input graph G. Let $r(\mathcal{A}, C, G) = \max_{v \in V(C)} r(v, \mathcal{A}, C, G)$. Then the *round complexity* of the algorithm $\mathcal{A}$ on n vertices is the maximum of $r(\mathcal{A}, C, G)$, taken over all n-vertex graphs C and G. This choice of metric is motivated by the fact that, in practice, the time for a message to traverse a network link is several orders of magnitude larger than local processor speed, so the running time is dominated by the time to send messages. This is also the reason why we can allow unlimited local computation in each round.

1.3 Problem Statements

In this thesis, we study the *subgraph detection* problem, defined as follows. Each node of C is given as input the same graph P, called the *pattern graph*, and the task is to determine if P is a subgraph of G. In particular, if P is a subgraph of G then we require at least one node to output TRUE, otherwise all nodes must output FALSE. Additionally, we consider the following variants of the subgraph detection problem. First, in the *induced subgraph detection* problem, we wish to detect whether or not P is an induced subgraph of G, rather than just a subgraph of G. Second, in the *vertex-rainbow subgraph detection* problem, each node of G has already been assigned a fixed colour and we wish to detect whether or not P appears as a subgraph of G under the additional constraint that the subgraph has at most one node of any colour. Similarly, we define the *edge-rainbow subgraph detection* problem, in which each edge of G has already been assigned a colour and we wish to detect whether or not P

appears as a subgraph of G under the additional constraint that the subgraph has at most one edge of any colour. Note that the vertex colouring and edge colouring of G are not necessarily proper colourings, i.e., it is possible that adjacent vertices or adjacent edges have been assigned the same colour. We may also combine two of these conditions to produce the *induced* {*vertex, edge*}*-rainbow subgraph detection* problem. We will also discuss the related *subgraph listing* problem, in which each node must output a list of subgraphs of G isomorphic to P so that each subgraph of G isomorphic to P appears in the output of at least one node.

Chapter 2

Related Work

2.1 Detecting a Cycle

2.1.1 Results for the CONGEST Model

Though the problem had previously been considered, much of the landscape was first discussed by Drucker et al. [11] in their seminal paper. They discuss lower bounds for the broadcast congested clique model, which is a variation of the CLIQUE model that requires nodes to send the same message to all other nodes in each round, and provide a framework using communication complexity techniques that had been adapted to show lower bounds in many other cases. It is also of interest that their lower bounds when P is a cycle also hold in the CONGEST model. First, if P has odd length greater than 3, they show that $\Omega(n/\log n)$ rounds are required. They also show that, if P is a cycle of length 4, then $\Omega(\sqrt{n}/\log n)$ rounds are required. Their work was later extended by Korhonen and Rybicki [19] to show that $\Omega(\sqrt{n}/\log n)$

rounds are required when P is a cycle of even length at least 6, in the CONGEST model. Their proof also shows that, if P is an even cycle of length at least 6 and G has degeneracy d, then there is no algorithm for finding P in the CONGEST model that uses $O(f(d)n^{\delta})$ rounds, where f is any function and $\delta < 1/2$. We discuss the lower bound proofs of Drucker et al. [11] and Korhonen and Rybicki [19] in more detail in Section 6.2. In terms of upper bounds, Korhonen and Rybicki [19] also show that when P is a path or a tree on k nodes then the problem can be solved in $O(k2^k)$ rounds, which they extend to show that if P is a cycle on k nodes or if P is a graph on k nodes containing exactly one cycle then the problem can be solved in $O(k2^kn)$ rounds. Fischer et al. [13] show that if P is a cycle of even length then the problem can be solved in $O(n^{1-1/(k(k-1))})$ rounds, which in particular beats the lower bound for odd cycles by Drucker et al. [11]. These results were improved on by Eden et al. [12], slightly improving the running time and showing that obtaining a better lower bound than Korhonen and Rybicki [19] for even cycles would show new lower bounds in circuit complexity, backing up the idea that the case when P is an odd cycle is somehow harder than the case when P is an even cycle. This is somewhat counter-intuitive, though it seems that this is caused by the fact that even cycles are perfect graphs whereas odd cycles are not.

2.1.2 Results for the CLIQUE Model

In the CLIQUE model, Censor-Hillel et al. [6] develop a technique to multiply matrices, showing that the case when P is a cycle of length 4 can be solved in a constant number of rounds and that the girth of G can be computed in $O(n^{0.158})$

rounds (where the girth is the length of the shortest cycle in G). These results would later be improved on by Censor-Hillel et al. [4], showing that, in the CLIQUE model, the girth of G can be computed in a constant number of rounds with a possible error of $+1$ and that the case when P is a cycle of even length can be solved in a constant number of rounds. These algorithms depend on some interesting work done by Lenzen and Wattenhofer [24], Patt-Shamir and Teplitsky [27], and Lenzen [23] who show that the classic problems of routing, sorting, and selection can be solved in a constant number of rounds in the CLIQUE model.

2.2 Detecting a Triangle

2.2.1 Results for the CONGEST Model

In the above discussion, the case when P is a cycle on three nodes (called a "triangle") has been omitted. As Drucker et al. [11] show, the standard communication complexity technique for lower bounds does not work in this case. Abboud et al. [1] and Fischer et al. [13] have attempted to prove lower bounds for the case when P is a triangle in the CONGEST model using a technique based on the pigeonhole principle, however their work has not yielded any non-trivial bounds (the best lower bound so far simplifies to $\Omega(1)$ rounds). In the CONGEST model, most work has been focused on the triangle listing problem. Izumi and Le Gall [18] provide algorithms that solve the detection problem when P is a triangle in $O(n^{2/3}(\log n)^{2/3})$ rounds and the listing problem when P is a triangle in $O(n^{3/4}\log n)$ rounds. This algorithm was improved on by Chang et al. [7] who use a randomized routing technique invovlving an expander decomposition to solve the triangle listing problem in $\tilde{O}(n^{1/2})$ rounds (the tilde means that there is a multiplicative polylogarithmic factor that is being omitted for the sake of conciseness). Though the details of this technique are complex, it is important to mention that expander decomposition depends on the routing techniques of Ghaffari et al. [16] and Ghaffari and Li [17], and enables a parameterization in terms of the mixing time of a subgraph of G. This algorithm would later be improved on by Chang and Saranurak [8] who modify the expander decomposition technique to produce an algorithm that solves triangle listing in $\tilde{O}(n^{1/3})$ rounds.

2.2.2 Results for the CLIQUE Model

An algorithm for the subgraph detection problem in the CLIQUE model was presented by Dolev et al. [10], parameterized in terms of t, the number of triangles in G. Their algorithm is randomized and has an expected round complexity of $O(n^{1/3}/(t^{t/3}+1))$. This was improved on by Censor-Hillel et al. [6] who show that the number of triangles in G can be counted in $O(n^{0.158})$ rounds using their matrix multiplication technique. Note that being able to count the number of triangles in G trivially solves the triangle detection problem, so this result is substantially stronger than that of Dolev et al. [10]. Izumi and Le Gall [18] show a lower bound of $\Omega(n^{1/3}/\log n)$ rounds for the triangle listing problem in the CLIQUE model using an information theoretic approach. This result is particularly significant since, as mentioned, no non-trivial lower bounds are known for the detection problem. Worth mentioning here is the work of Pandurangan et al. [26] who show a lower bound on triangle listing in the k machines model, which is similar to the CLIQUE model but is designed to better capture parallelization. What is relevant here is the fact that their result implies the same lower bound as the one shown by Izumi and Le Gall [18]. Notice that this lower bound is only a polylogarithmic factor away from the algorithm of Chang and Saranurak [8]. This is interesting as this algorithm and lower bound all hold in both the CONGEST model and the CLIQUE model, suggesting that, unlike in the case of triangle detection, the extra power of the CLIQUE model may not help for listing triangles.

2.3 Detecting a Clique

2.3.1 Results for the CONGEST Model

Notice that a triangle is also a complete graph (called a "clique") on three nodes, and the detection and listing problems have also been studied for larger cliques. A representative example is the work of Brakerski and Patt-Shamir [2] who consider the case when P is a near-clique (a near-clique is a graph formed by taking a clique and removing a constant fraction of its edges) using property testing, a randomized technique, and constant time in the CONGEST model. However, recent research has seen significant results in adapting the triangle listing techniques above to work for any clique. Intuitively this makes some sense as the naive approach to detecting cliques is to have each node learn the neighbours of its neighbours, which will find all the cliques in G of any size. First, if P is a clique of size at least four, then the communication complexity lower bound technique will work here. It was applied by Czumaj and Konrad [9] to show a lower bound of $\Omega\left(\frac{n}{(l+\sqrt{n})\log n}\right)$ rounds for the detection problem if P is a clique of size l and $l \geq 4$, in the CONGEST model. They also prove that the communication complexity technique cannot give a better lower bound in this case. On the algorithms side, Eden et al. [12] show that cliques on four and five nodes can be listed in $O(n^{5/6+o(1)})$ and $O(n^{21/22+o(1)})$ rounds, respectively, in the CONGEST model. This was expanded on by Censor-Hillel et al. [5] and Censor-Hillel et al. [3] using the expander decomposition technique, culminating in an algorithm that lists all cliques of size l in G in $\tilde{O}(n^{1-2/l})$ rounds in the CONGEST model. It is interesting to note that this differs from the detection lower bound of Czumaj and Konrad [9] by only a polygarithmic factor, suggesting that this algorithm is close to the best we can do for the detection problem.

2.3.2 Results for the CLIQUE Model

In the CLIQUE model, Fischer et al. [13] extended the work of Izumi and Le Gall [18] to show a lower bound of $\tilde{\Omega}(n^{1-2/l})$ for listing cliques of size l. Notice that this lower bound is only a polylogarithmic factor away from the algorithm of Censor-Hillel et al. [3]. These results show that the interesting behaviour that we observed when P was a triangle also hold when P is a clique of arbitrary size; in particular that the round complexity of listing cliques is perhaps the same in the CONGEST model and CLIQUE model. Intuitively, we can see that this is because the additional power that the CLIQUE model has over the CONGEST model is the ability to quickly learn information that is non-local in C, however clique listing only requires learning purely local information.

2.4 Arbitrary Pattern Graphs

2.4.1 Results for the CONGEST Model

Another interesting avenue of research is to design algorithms that will work for *any* pattern graph P. In the CONGEST model, Fischer et al. [13] show that there is a graph P with $O(k)$ vertices that requires $\Omega(n^{2-1/k}/(k \log n))$ rounds to detect, extending previous work by Fischer et al. [14]. In particular, this implies an almost quadratic lower bound if P can be any graph. On the upper bound side, Eden et al. [12] show that, for any graph P with k vertices, there is an algorithm in the CONGEST model that detects P in $\tilde{O}(n^{2-\frac{2}{3k+1}+o(1)})$ rounds, showing that the lower bound proved by Fischer et al. [13] is close to tight.

2.4.2 Results for the CLIQUE Model

Dolev et al. [10] show that, in the CLIQUE model, if P has d nodes then there is an algorithm that detects P in $O\left(n^{(d-2)/d}/\log n\right)$ rounds. Further, they show that if G has maximum degree Δ and P has diameter D (the *diameter* of a graph is the maximum of the lengths of the shortest paths between two vertices of the graph) then there is an algorithm that detects P in $O(\Delta^{D+1}/n)$ rounds in the CLIQUE model.

2.5 Induced Subgraphs and Rainbow Variants

2.5.1 Results for the CONGEST Model

Korhonen and Nikabadi [25] show that detecting induced P when P is an even cycle in the CONGEST model requires at least $\Omega(n/\log n)$ rounds. Further, they show that detecting induced P takes at least $\Omega(n^{1/2}/\log n)$ rounds if P contains a 4-clique and that there is a graph P so that P has treewidth 2 and takes almost quadratic time to detect. Similarly, they show that detecting a vertex-rainbow cycle takes at least $\Omega(n/\log n)$ rounds and that detecting a vertex-rainbow induced path of length k takes at least $\Omega(n/\log n)$ rounds. They then show that bounding the degeneracy of G does not improve the almost quadratic lower bound for some P of treewidth 2, though it does improve the running time for algorithms that detect induced cycles and induced trees. Le Gall and Miyamoto [22] expand on this, showing that detecting induced P, when P has k nodes, can be done in $\tilde{O}(n^{2-2/(3k+1)})$ rounds, for any P, and that detecting an induced k-cycle takes at least $\Omega(n/\log n)$ rounds. Of particular interest is the fact that this lower bound, as well as the $\Omega(n/\log n)$ lower bound for rainbow cycle detection, does not depend on the parity of k, showing that the difference between even and odd cycles does not seem to be present in these variants. Finally, we comment that, to the best of our knowledge, our work is the first to introduce the edge-rainbow subgraph detection problem.

Chapter 3

Our Results

Our work is based on the following intuition: If P and Q are graphs so that Q is an induced subgraph of P, then detecting P should be at least as hard as detecting Q. Notice that, if we do not require that the subgraph be induced, then this is not true. To see this, consider C_5 and K_5. Certainly, C_5 is a subgraph of K_5. However, as we have discussed, detecting C_5 in CONGEST takes at least $\Omega(n/\log n)$ rounds [11] and K_5 can be detected in $\tilde{O}(n^{3/5})$ rounds in CONGEST [3]. It is not clear to us how to formally prove that this intuition is correct in general, so we restrict ourselves to the case in which Q is a cycle. In Chapter 4, we describe the general technique that we use to prove all of our results. In Chapter 5, by extending the work of Korhonen and Nikabadi [25], we prove that if P is a connected graph that contains an induced cycle, then detecting a vertex-rainbow or edge-rainbow copy of P takes at least $\Omega(n/\log n)$ rounds in the CONGEST model. We observe that, based on this result, it seems possible that a similar result could hold for the general/uncoloured subgraph detection problem. To explore this idea, in Chapter 6, we consider the case of detecting P when P is a necklace graph. This is motivated by the fact that necklace graphs are a simple, yet non-trivial, class of graphs that we know each contain induced cycles. In particular, we show that, if P is a necklace graph that contains an odd cycle of length at least 5, then detecting a copy of P takes at least $\Omega(n/\log n)$ rounds in the CONGEST model. By doing so, we show that it may be possible to get more general results, closer to those we show in the rainbow case, and highlight the challenges in adapting the proof to the uncoloured case. After presenting the proof, we discuss why we are unable to use this technique to derive lower bounds in the cases in which P contains a triangle or only even cycles.

Chapter 4

Lower Bound Framework

For our results, we use reductions from 2-party communication complexity, so we can make use of known lower bounds in this model. This is the most common technique used for proving lower bounds for subgraph detection in congested networks, as seen in, e.g., Drucker et al. [11], Korhonen and Nikabadi [25], Le Gall and Miyamoto [22], and Czumaj and Konrad [9]. In the *2-party communication complexity* model, we assume that there are two parties, usually called "Alice" and "Bob", who are each given a bit string of the same length. Say that Alice is given the bit string Ξ and Bob is given the bit string Ψ so that Ξ and Ψ each have length L. As in the models we described in Section 1.2, we assume that time proceeds in discrete, synchronous rounds. Also as in those models, each round consists of three phases. First, Alice and Bob may each perform unlimited local computation. Then, exactly one of Alice or Bob may send exactly one bit to the other. Finally, the bit that was sent is delivered to the other party. To determine which of them will send a bit, Alice and Bob use a previously-defined protocol which also describes how they can use their input and any messages they may have received before this point to determine which bit to send. Note that this implicitly assumes that Alice and Bob have unique ID's, i.e., Alice knows she is "Alice" and Bob knows he is "Bob". Now, Alice and Bob must terminate their execution after a finite number of rounds and specify an output bit. Let Γ be the set of all bit strings of length L. Let $f : \Gamma \times \Gamma \to \{0, 1\}$ be a Boolean function. We will often call such a function f a *problem* for this model. For example, the equality problem EQ is defined as $f(\Xi, \Psi) = 1$ if and only if $\Xi = \Psi$ and the set disjointness problem DISJ is defined as $f(\Xi, \Psi) = 0$ if and only if there is at least one $p \in |\Xi|$ such that $\Xi_p = \Psi_p = 1$. A protocol is said to *solve* f if, for every pair $\Xi, \Psi \in \Gamma$, Alice and Bob each produce the result $f(\Xi, \Psi)$ upon terminating while following the protocol. Let $f : \Gamma \times \Gamma \to \{0, 1\}$ be a function and let $\mathcal{P}$ be a protocol that solves f. Then the *communication complexity of $\mathcal{P}$ for f on length L* is the maximum number of rounds that it takes for Alice and Bob to both terminate when running the protocol $\mathcal{P}$ on bit strings Ξ and Ψ, taken over all $\Xi, \Psi \in \Gamma$. The *communication complexity* of the function f on length L is defined as the minimum communication complexity of any protocol that solves f. Note that, in a naive manner, any function f has communication complexity at most $L + 1$. To see this, notice that Alice can spend the first L rounds sending each bit of Ξ to Bob. Then, Bob can simply compute $f(\Xi, \Psi)$ and send the result to Alice, at which point they can each terminate and produce $f(\Xi, \Psi)$, taking $L + 1$ rounds.

Now, to reduce from a problem f to the subgraph detection problem in the CONGEST model, Alice and Bob construct a graph G, which depends on Ξ, Ψ, and the pattern graph P, and simulate a CONGEST algorithm running on G. For the sake of common terminology, such a graph G is called a *lower bound graph*. Let $\mathcal{A}$ be any algorithm that detects P in the CONGEST model. We partition the vertices of G into two sets A and B, so that Alice simulates the execution of algorithm $\mathcal{A}$ by each node in A and Bob simulates the execution of algorithm $\mathcal{A}$ by each node in B. There are two key properties that our construction of G should satisfy. First, Alice and Bob should be able to construct G without using any communication, so that every bit that Alice sends to Bob is part of a message from a vertex in A to a vertex in B and every bit that Bob sends to Alice is part of a message from a vertex in B to a vertex in A (with one exception that we will mention shortly). In particular, this means that only $G[A]$ should depend on Ξ and only $G[B]$ should depend on Ψ. Second, G should contain a copy of P if and only if $f(\Xi, \Psi) = 0$. Then, after each node has terminated, Alice will look at the output of all the A nodes and Bob will look at the output of all the B nodes. They will then send the other a '1' if at least one node outputs 1 and a '0' otherwise. Thus, $f(\Xi, \Psi) = 1$ if and only if both sent bits are 0 and Alice and Bob can output accordingly. Now, suppose that there is a known lower bound of R rounds for solving f in the 2-party communication model. Then, since we have described a protocol to solve f, at least R bits are exchanged between Alice and Bob. This means that one of them, say Alice, sends the other at least $R/2$ bits. Now suppose that there are C edges in G with one endpoint in A and the other in B. Such edges are called *cut edges*. Then, by the Pigeonhole Principle, at least one node a in A sends a node b in B at least $\frac{R-1}{2C}$ bits during the simulation (the "-1" comes from the bit used to share the result). Finally, since each node can send at most $\log n$ bits in each round (ignoring asymptotic notation for simplicity),

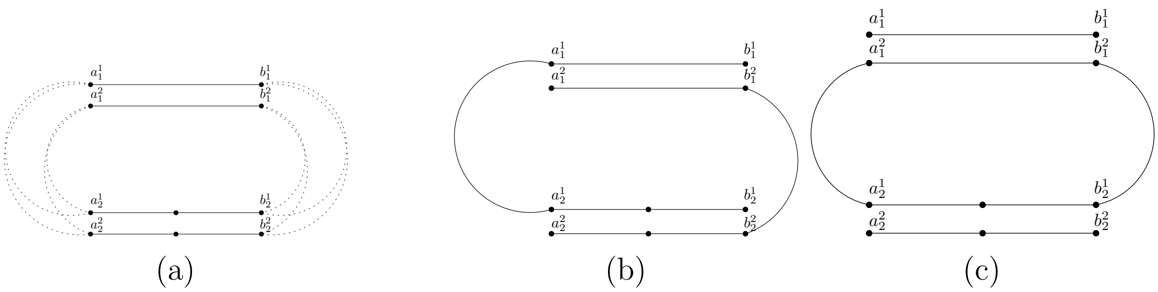

Figure 4.1: (a) A lower bound graph for C_5. The dotted edges represent edges that could be added by Alice and Bob depending on the values of Ξ and Ψ. (b) A lower bound graph for C_5 where Alice and Bob have each added one edge, but no C_5 is present. Note that they do not have any bit set to '1' in common. (c) A lower bound graph for C_5 where Alice and Bob have each added one edge and C_5 is present. Note that the bit they have set to '1' is in the same position for each of them.

this means that a executes for at least $\frac{R-1}{2C\log n}$ rounds. Thus, the algorithm $\mathcal{A}$ takes at least $\frac{R-1}{2C\log n}$ rounds and, since $\mathcal{A}$ was arbitrary, this means that any algorithm that detects P in the CONGEST model takes at least $\frac{R-1}{2C\log n}$ rounds for some input graph G, implying a lower bound of $\frac{R-1}{2C\log n}$ on detecting P in the CONGEST model.

For our reductions, we will take the function f to be the set disjointness problem, i.e., $f(\Xi, \Psi) = 0$ if and only if there is at least one $p \in [L]$ such that $\Xi_p = \Psi_p = 1$. It is known that the communication complexity of the set disjointness problem is $L+1$, so we may set $R = L + 1$ in the argument described above. A proof of this can be found in Kushilevitz and Nisan [20].

As an example of applying this technique, we briefly summarize the proof of Drucker et al. [11] that, if $l \geq 5$ is a fixed odd integer constant, then detecting C_l takes at least $\Omega(n/\log n)$ rounds in the CONGEST model. Let N be an arbitrary positive integer and consider an instance of set disjointness with input strings of length $L = N^2$. Let Ξ and Ψ be the inputs to Alice and Bob, respectively. Now, we describe the graph G that Alice and Bob will construct. In Figure 4.1, we show an example of this when $l = 5$ and $N = 2$. First, take N paths with $\lfloor l/2 \rfloor$ vertices (in the figure, 2 vertices). For each $p \in [N]$, label the leftmost vertex of one of these paths a_1^p and the rightmost vertex b_1^p. Next, take N paths with $\lceil l/2 \rceil$ vertices (in the figure, 3 vertices). For each $p \in [N]$, label the leftmost vertex of one of these paths a_2^p and the rightmost vertex b_2^p. For each $p \in [N]$ and each $q \in [2]$, Bob will simulate the vertex b_q^p, i.e., $B = \{b_q^p : p \in [N], q \in [2]\}$. Alice will simulate all remaining vertices, i.e., $A = V(G) \setminus B$. It follows that the number of cut edges, i.e., the number of edges with exactly one endpoint in A and one endpoint in B, is $2N$. Notice that, up to this point, Alice and Bob both know the entire construction (i.e., they both know the graph illustrated in Figure 4.1a). Now, fix a bijection $g : [N^2] \to [N] \times [N]$. For each $p \in [N^2]$ so that $\Xi_p = 1$, write $g(p) = (p_1, p_2)$. Then Alice will add an edge between the vertices $a_1^{p_1}$ and $a_2^{p_2}$. Similarly, for each $p \in [N^2]$ so that $\Psi_p = 1$, write $g(p) = (p_1, p_2)$. Then Bob will add an edge between the vertices $b_1^{p_1}$ and $b_2^{p_2}$. The key property of this construction is the following.

Lemma 4.1 (Lemma 18 in Drucker et al. [11])**.** *The graph G (described above) contains a copy of C_l if and only if there is a $p \in [N^2]$ such that $a_1^{p_1} a_2^{p_2} \in E(G)$ and $b_1^{p_1} b_2^{p_2} \in E(G)$, where $g(p) = (p_1, p_2)$.*

Note that Lemma 4.1 is equivalent to the property that G contains a copy of C_l if and only if there is a $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$. Thus, since Alice and Bob constructed this graph without using any communication, the graph G satisfies the properties that we require. Then, it remains to notice that the number of vertices is $|V(G)| = n = N \cdot l \in \Theta(N)$ and the number of cut edges in G is $C = 2N$ to get the lower bound of

$$\frac{N^2+1-1}{2 \cdot 2N \log n} = \frac{N}{4 \log n} \in \Omega\left(\frac{n}{\log n}\right)$$

as desired.

Chapter 5

Rainbow Subgraphs

5.1 Lower Bounds

The goal of this section is to prove the following lower bounds on the round complexity of any distributed rainbow subgraph detection algorithm when the pattern graph contains an induced cycle of length at least 4.

Theorem 5.1. *Let P be a fixed graph that contains an induced cycle of length at least 4. Consider the vertex-rainbow subgraph detection problem in the CONGEST model on all graphs G with n vertices. Then any algorithm that solves the problem takes at least $\Omega(n/\log n)$ rounds in the worst case.*

Theorem 5.2. *Let P be a fixed, connected graph that contains an induced cycle of length at least 4. Consider the edge-rainbow subgraph detection problem in the CONGEST model on all graphs G with n vertices. Then any algorithm that solves the problem takes at least $\Omega(n/\log n)$ rounds in the worst case.*

Large parts of the proofs of these statements are agnostic to the type of colouring (i.e., vertex-rainbow versus edge-rainbow), so we will present these parts once in Section 5.1.1 and defer the parts of the proofs that depend on the colourings in Sections 5.1.2 and 5.1.3.

5.1.1 The Generic Construction

Consider any graph P with k vertices that contains an induced cycle C of length $l \geq 4$. Let a_1a_2 and b_1b_2 be disjoint edges in the induced cycle so that $|\{a_1, a_2, b_1, b_2\}| = 4$. For some $l_1, l_2 \geq 0$, write the full induced cycle as $a_1, u_1, u_2, \ldots, u_{l_1}, b_1, b_2, v_{l_2}, v_{l_2-1}, \ldots, v_1, a_2, a_1$. Note that l_1 or l_2 may be equal to 0 (or both of them may be) and that $l = l_1 + l_2 + 4$. Finally, let $w_1, w_2, \ldots, w_{k-l}$ be the remaining vertices of P. For convenience, we will write $W = \{w_1, w_2, \ldots, w_{k-l}\}$ and $C = \{a_1, u_1, u_2, \ldots, u_{l_1}, b_1, b_2, v_{l_2}, v_{l_2-1}, \ldots, v_1, a_2\}$, so that W and C partition $V(P)$.

Let $N > 1$ be any positive integer and consider an instance of the standard 2-party set disjointness problem of size N^2. Let Ξ be the input to Alice and let Ψ be the input to Bob. Recall that each of these is a bit string of length N^2. We will show that Alice and Bob can simulate any algorithm to detect rainbow P and use it to solve this instance. To do this, they will construct a graph G depending on Ξ and Ψ that represents a CONGEST network, each of them simulating some of the nodes in the network. We will now describe how to construct the graph G, and we suggest that the reader consult Figure 5.1. First, consider the path $a_1, u_1, \ldots, u_{l_1}, b_1$ in the induced cycle C. We make N copies of this path, i.e., for each $p \in [N]$, we create a path $a_1^p, u_1^p, \ldots, u_{l_1}^p, b_1^p$. Similarly, we create N copies of the path $a_2, v_1, \ldots, v_{l_2}, b_2$,

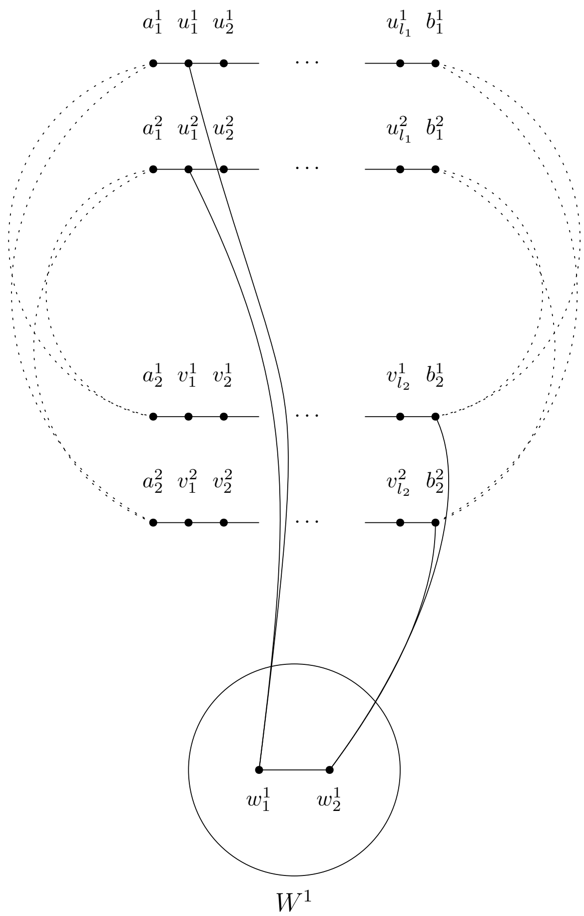

Figure 5.1: An example of constructing the graph G with $N = 2$. The dotted edges represent all the edges that could be added depending on the values of Ξ and Ψ.

that is, for each $p \in [N]$, create a path $a_2^p, v_1^p, \ldots, v_{l_2}^p, b_2^p$. Next, add a copy of W, i.e., the remaining vertices in P that aren't part of the part of the induced cycle C. In particular, add vertices $w_1^1, \ldots, w_{k-l}^1$, and denote this set of copies as W^1. Then, add an edge between the pair $w_{q_1}^1$ and $w_{q_2}^1$ if and only if w_{q_1} and w_{q_2} are adjacent in P. Finally, we add edges from the copies of vertices of C to the copy of vertices from W. For $p \in [N]$ and $q \in [k-l]$, add the edge $a_1^p w_q^1$ if a_1 and w_q are adjacent in P, the edge $a_2^p w_q^1$ if a_2 and w_q are adjacent in P, the edge $b_1^p w_q^1$ if b_1 and w_q are adjacent in P, and the edge $b_2^p w_q^1$ if b_2 and w_q are adjacent in P. Similarly, for

$p \in [N], q_1 \in [l_1]$, and $q_2 \in [k-l]$, add the edge $u^p_{q_1} w^1_{q_2}$ if u_{q_1} and w_{q_2} are adjacent in P and for $p \in [N], q_1 \in [l_2]$, and $q_2 \in [k-l]$ add the edge $v^p_{q_1} w^1_{q_2}$ if v_{q_1} and w_{q_2} are adjacent in P. Now, fix an arbitrary bijection $f : [N^2] \to [N] \times [N]$. To slightly abuse notation, write that $f(p) = (p_1, p_2)$ for each $p \in [N^2]$. Now for each $p \in [N^2]$, if Ξ_p is 1, Alice will add the edge $a^{p_1}_1 a^{p_2}_2$ to G and if Ψ_p is 1, Bob will add the edge $b^{p_1}_1 b^{p_2}_2$ to G.

In Sections 5.1.2 and 5.1.3, we describe how to colour G and we prove that G contains a rainbow copy of P if and only if there is a $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$.

Next, suppose we have any distributed algorithm $\mathcal{A}$ that detects a rainbow copy of P in G. Alice and Bob will solve set disjointness in the 2-party communication complexity model as follows. Let $B = \{b^p_q : p \in [N], q \in [2]\} \subset V(G)$ and let $A = V(G) \setminus B \subset V(G)$. Alice simulates the algorithm $\mathcal{A}$ on the nodes in A and Bob simulates the algorithm $\mathcal{A}$ on the nodes in B, sending each other the messages that pass over the cut A, B. When algorithm $\mathcal{A}$ outputs its answer, i.e., whether or not G contains a rainbow copy of P, Alice and Bob can deduce correct answer of the set-disjointness instance since G contains a rainbow copy of P if and only if there is a $p \in [N^2]$ such that $\Xi_p = \Psi_p = 1$. Therefore, by known lower bounds for the set disjointness problem in the 2-party communication complexity model, there exists an instance that requires one of Alice or Bob to send the other at least $\Omega(N^2)$ bits to do this. Since they only send each other the messages that pass over the cut A, B when simulating the algorithm, at least $\Omega(N^2)$ bits of communication pass over the cut A, B in one direction throughout the execution of the algorithm. Now consider the neighbours of any vertex $b^p_q \in B$ that are in A. There is exactly one vertex from the path connecting a_q^p and b_q^p and at most $|W| \leq k \in O(1)$ in W^1. Thus there are a constant number of such neighbours, so the size of the cut is $O(N)$ (since $|B| = 2N$). Therefore there is a node in A and a node in B so that one of them sends the other at least $\Omega(N)$ bits of communication throughout the execution of the algorithm. Then since the bandwidth per edge is assumed to be $O(\log n)$, this takes at least $\Omega(N/\log n)$ rounds, where $n = |V(G)|$. The constructed graph consists of N copies of the cycle C with length l, and there are $k - l$ additional vertices in P outside of C, so in total we have that the total number of vertices is $n = N \cdot l + k - l \in \Theta(N)$. Therefore, we get that the algorithm takes at least $\Omega(n/\log n)$ rounds, and since this was true for any algorithm, this shows that any algorithm that detects vertex-rainbow or edge-rainbow P in the CONGEST model takes at least $\Omega(n/\log n)$ rounds, as claimed.

5.1.2 The Vertex-Rainbow Proof

First, we colour the vertices of G. Intuitively, each vertex of P is associated with a unique colour and each "copy" of that vertex in G is assigned that colour. For each $p \in [N]$, assign the vertex a_1^p the colour 1, the vertex a_2^p the colour 2, the vertex b_1^p the colour 3, and the vertex b_2^p the colour 4. For each $p \in [N]$ and $q \in [l_1]$, assign the vertex u_q^p the colour $4 + q$. For each $p \in [N]$ and $q \in [l_2]$, assign the vertex v_q^p the colour $4 + l_1 + q$. Finally, for each $q \in [k - l]$, assign the vertex w_q^1 the colour $4 + l_1 + l_2 + q$. Thus, we have coloured all the vertices of G. Note that we have used exactly $4 + l_1 + l_2 + (k - l) = l + (k - l) = k$ colours.

Now we will show that G contains a vertex-rainbow copy of P if and only if there is a $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$.

First, suppose that there is an $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$ and let $(p_1, p_2) = f(p)$. Then, Alice will add the edge $a_1^{p_1} a_2^{p_2}$ and Bob will add the edge $b_1^{p_1} b_2^{p_2}$, so the subgraph of G induced by $a_1^{p_1}, u_1^{p_1}, \ldots, u_{l_1}^{p_1}, b_1^{p_1}, b_2^{p_2}, v_{l_2}^{p_2}, \ldots, v_1^{p_2}, a_2^{p_2}, w_1^1, \ldots, w_{k-l}^1$ is a vertex-rainbow copy of P since the added edges create a copy of the cycle C.

Conversely, suppose that G contains a vertex-rainbow copy of P. We will show that there is a $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$. A vertex-rainbow copy of P has k vertices with k unique colours, so from our colouring of G, the vertex-rainbow copy of P must consist of one copy of each vertex $a_1, a_2, b_1, b_2, u_1, \ldots, u_{l_1}, v_1, \ldots, v_{l_2}, w_1, \ldots, w_{k-l}$. Since G only contains one copy of each w_q, it follows that the vertex-rainbow copy of P contains the vertices $W^1 = \{w_1^1, \ldots, w_{k-l}^1\}$, which corresponds to the set W in P. The remaining vertices in the vertex-rainbow copy of P form a cycle, and we refer to these vertices as the set C^*, which corresponds to the cycle C in P. Since there are N copies of each vertex of C in G, there exist $p_1, \ldots, p_{l_1+l_2+4}$ such that $C^* = \{a_1^{p_1}, u_1^{p_2}, \ldots, u_{l_1}^{p_{l_1+1}}, b_1^{p_{l_1+2}}, a_2^{p_{l_1+3}}, v_1^{p_{l_1+4}}, \ldots, v_{l_2}^{p_{l_1+l_2+3}}, b_2^{p_{l_1+l_2+4}}\}$. Let $S = W^1 \cup C^*$. Let $P_1^* = \{a_1^{p_1}, u_1^{p_2}, \ldots, u_{l_1}^{p_{l_1+1}}, b_1^{p_{l_1+2}}\}$ and $P_2^* = \{a_2^{p_{l_1+3}}, v_1^{p_{l_1+4}}, \ldots, v_{l_2}^{p_{l_1+l_2+3}}, b_2^{p_{l_1+l_2+4}}\}$ so that C^* is partitioned by P_1^* and P_2^*. Now, since there is a copy of P in G with vertex set S, the subgraph of G induced by S has at least as many edges as P, i.e., $|E(P)| \leq |E(G[S])|$. Let E_1 be the set of edges of P with at least one endpoint in W and let E_2 be the set of edges of $G[S]$ with at least one endpoint in W^1. Define a bijection $g : E_1 \to E_2$ by $w_{q_1} w_{q_2} \mapsto w_{q_1}^1 w_{q_2}^1$, if $w_{q_1}, w_{q_2} \in W$ and $wz \mapsto w^1 z^p$, for an index p, if $w \in W$ and $z \in C$. Note that g is well-defined since, for each $z \in C$, there is exactly one index p so that $z^p \in C^*$. Then since g is a bijection, $|E_1| = |E_2|$, and so $|E(G[S]) \setminus E_2| \geq |E(P) \setminus E_1| = |E(C)| = l$. In other words, the number of edges with no endpoints in W^1 (i.e. the number of edges with both endpoints in C^*) is at least l. We claim that this implies that $p_1 = p_2 = \cdots = p_{l_1+2}$ and $p_{l_1+3} = p_{l_1+4} = \cdots = p_{l_1+l_2+4}$, and $a_1^{p_1} a_2^{p_{l_1+3}}, b_1^{p_{l_1+2}} b_2^{p_{l_1+l_2+4}} \in E(G)$. To see this, consider all possible choices of $p_1, \ldots, p_{l_1+2}$ and count the number of edges in $G[P_1^*]$. By the construction of G, if we consider any internal node in the p'th copy of the path $a_1, u_1, \ldots, u_{l_1} b_1$, we see that both of its incident edges only involve nodes in the same copy of the path. Thus, there is a maximum of $l_1 + 1$ edges that occurs exactly when $p_1 = p_2 = \cdots = p_{l_1+2}$. Similarly, consider all possible choices of $p_{l_1+3}, \ldots, p_{l_1+l_2+4}$ and count the number of edges in $G[P_2^*]$. By the construction of G, if we consider any internal node in the p'th copy of the path $a_2, v_1, \ldots, v_{l_2}, b_2$, we see that both of its incident edges only involve nodes in the same copy of the path. Thus, there is a maximum of $l_2 + 1$ edges that occurs exactly when $p_{l_1+3} = p_{l_1+4} = \cdots = p_{l_1+l_2+4}$. So, we have shown that $|E(G[P_1*])| + |E(G[P_2^*])| \leq l_1 + l_2 + 2$. Next, we compute $|E(G[C^*])|$. The only possible edges we have not counted in this subgraph are those induced by the path endpoints $a_1^{p_1}$, $a_2^{p_{l_1+3}}$, $b_1^{p_{l_1+2}}$, $b_2^{p_{l_1+l_2+4}}$. By the construction of G, there are exactly two possible additional edges: $a_1^{p_1} a_2^{p_{l_1+3}}$ and $b_1^{p_{l_1+2}} b_2^{p_{l_1+l_2+4}}$. So, since P_1^* and P_2^* partition C^*, we have shown that there are at most $l_1 + l_2 + 4$ edges in the subgraph of G induced by C^*, and that there is exactly one way to achieve the maximum $l_1 + l_2 + 4$. As the number of edges in C is $l = l_1 + l_2 + 4$, it must be the case that the copy of C in G, i.e., the subgraph of G induced by C^*, must have exactly $l_1 + l_2 + 4$ edges. So, we conclude that $p_{l_1+3} = p_{l_1+4} = \cdots = p_{l_1+l_2+4}$, that $p_{l_1+3} = p_{l_1+4} = \cdots = p_{l_1+l_2+4}$, and that both of the edges $a_1^{p_1} a_2^{p_{l_1+3}}$ and $b_1^{p_{l_1+2}} b_2^{p_{l_1+l_2+4}} = b_1^{p_1} b_2^{p_{l_1+3}}$ are present. Letting $p = f^{-1}(p_1, p_{l_1+3})$, it follows from the construction of G that $\Xi_p = \Psi_p = 1$, as claimed.

5.1.3 The Edge-Rainbow Proof

First, we colour the edges of G. Let E_1 be the set of edges of P with both endpoints in C, E_2 be the set of edges of P with exactly one endpoint in C, and E_3 be the set of edges of P with no endpoints in C.

First, we colour the edges in G that are copies of edges in E_1. In particular, these are the edges in the N paths with with endpoints a_1^p, b_1^p and endpoints a_2^p, b_2^p, as well as any edges between endpoints of different paths. For each pair $p_1, p_2 \in [N]$, if the edge $a_1^{p_1} a_2^{p_2}$ is present in G then assign it the colour 1. Similarly, for each pair $p_1, p_2 \in [N]$, if the edge $b_1^{p_1} b_2^{p_2}$ is present in G then assign it the colour 2. Next, for each $p \in [N]$, assign the edge $a_1^p u_1^p$ the colour 3, assign the edge $u_{l_1}^p b_1^p$ the colour 4, the edge $a_2^p v_1^p$ the colour 5, and the edge $v_{l_2}^p b_2^p$ the colour 6. Then for each $i \in [N]$ and each $q \in [l_1 - 1]$, assign the edge $u_q^p u_{q+1}^p$ the colour $6 + q$. Similarly, for each $p \in [N]$ and each $q \in [l_2 - 1]$, assign the edge $v_q^p v_{q+1}^p$ the colour $6 + l_1 - 1 + q$. At this point, the colours $1, \ldots, l$ have been used to colour the edges of E_1.

Next, we colour the edges of G that are copies of the edges in E_2. We denote the copy of an edge $e \in E_2$ using the notation e^2 and we denote the copy of an edge $e \in E_3$ using the notation e^3. Arbitrarily order the edges of E_2, say $E_2 = \{e_1^2, \ldots, e_{|E_2|}^2\}$. Informally, each edge $e_q \in E_2$ will be associated with its own colour $l + q$. Each such edge e_q has one endpoint $w \in W$ and one endpoint $z \in C$, and for all the edges in G between the copy of w in G and each copy of z in G, we associate them to e_q by assigning them the colour $l + q$. More formally, for each $q \in [|E_2|]$, let $e_q^2 = wz$, where $w \in W$ and $z \in C$, and for each $p \in [N]$, assign the edge $w^1 z^p$ of G the colour $l + q$. Similarly, arbitrarily order the edges of E_3, say $E_3 = \{e_1^3, \ldots, e_{|E_3|}^3\}$. Informally, each edge in E_3 has exactly one copy in G, and we assign a unique colour to each. More formally, for each $q \in [|E_3|]$, let $e^3_q = w_1 w_2$, where $w_1, w_2 \in W$, and assign the edge $w^1_1 w^1_2$ of G the colour $l + |E_1| + q$. This completes the colouring of all edges in G. Intuitively, each edge of P is associated with a unique colour and each "copy" of that edge in G is assigned that colour, so we have used exactly $|E(P)|$ colours.

Now we will show that G contains an edge-rainbow copy of P if and only if there is a $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$. First, suppose that there is a $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$. Then, Alice will add the edge $a^{p_1}_1 a^{p_2}_2$ and Bob will add the edge $b^{p_1}_1 b^{p_2}_2$, so the subgraph of G induced by $a^{p_1}_1, u^{p_1}_1, \ldots, u^{p_1}_{l_1}, b^{p_1}_1, b^{p_2}_2, v^{p_2}_{l_2}, \ldots, v^{p_2}_1, a^{p_2}_2, w^1_1, \ldots, w^1_{k-l}$ is an edge-rainbow copy of P since the added edges create a copy of the cycle C.

Conversely, suppose that G contains an edge-rainbow copy of P. We will show that there is a $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$. Let S be the vertex set of the edge-rainbow copy of P and let $W^1 = \{w^1_1, \ldots, w^1_{k-l}\}$. Now, an edge-rainbow copy of P has $|E(P)|$ edges with $|E(P)|$ unique colours. Since our colouring of G uses $|E(P)|$ colours, this implies that the edge-rainbow copy of P contains an edge of each colour that we assigned. So, in particular, since the copy of each edge in E_3 was assigned a unique colour, we conclude that the edge-rainbow copy of P in G contains exactly one edge $w^1_1 w^1_2$ for each edge $w_1 w_2 \in E_3$. Similarly, for each edge $wz \in E_2$, we associated a unique colour and coloured all copies of it in G with that colour, so it follows that the edge-rainbow copy of P in G contains exactly one edge $w^1 z^q$, for some index $q \in [N]$. Thus, the edge-rainbow copy of P contains all the vertices of W^1 (since P is connected, so there are no isolated vertices of P). Thus, there are exactly $|P| - |W| = |C| = l$ vertices in the edge-rainbow copy of P in $V(G) \setminus W^1$.

Now consider the edges of E_1, which form the cycle $P[C]$. Recall that our colouring of G uses exactly $|E(P)|$ colours, so the edge-rainbow copy of P in G contains one edge that was coloured c for each $c \in \{1, \dots, l\}$. From the definition of our colouring of G, it follows that the edge-rainbow copy of P in G contains the following l edges, for some choice of superscripts, each in the range $[N]$: first, there is exactly one edge $a_1^{p_1} a_2^{p_2}$ and exactly one edge $b_1^{p_3} b_2^{p_4}$. Then there is exactly one edge $a_1^{p_{a_1}} u_1^{p_{a_1}}$, exactly one edge $u_{l_1}^{p_{u_{l_1}}} b_1^{p_{u_{l_1}}}$, exactly one edge $a_2^{p_{a_2}} v_1^{p_{a_2}}$, and exactly one edge $v_{l_2}^{p_{v_{l_2}}} b_2^{p_{v_{l_2}}}$. Finally, for each $q \in [l_1 - 1]$ there is exactly one edge $u_q^{p_{u_q}} u_{q+1}^{p_{u_q}}$ and for each $q \in [l_2 - 1]$ there is exactly one edge $v_q^{p_{v_q}} v_{q+1}^{p_{v_q}}$. Now, consider all combinations of values for the superscripts $p_1, p_2, p_3, p_4, p_{a_1}, p_{u_{l_1}}, p_{a_2}, p_{v_{l_2}}, p_{u_q}$ for each $q \in [l_1 - 1]$, p_{v_q} for each $q \in [l_2 - 1]$. Each combination specifies a set of edges in G and, for each combination, we count the number of vertices in the subgraph consisting of these edges. The minimum possible number of vertices in such a subgraph is equal to the number of edges l, and this occurs exactly when the subgraph is a cycle, i.e., when $p_1 = p_{a_1} = p_{u_1} = \cdots = p_{u_{l_1-1}} = p_{u_{l_1}} = p_3$ and $p_2 = p_{a_2} = p_{v_1} = \cdots = p_{v_{l_2-1}} = p_{v_{l_2}} = p_4$. However, we saw in the previous paragraph that there are exactly l vertices in the edge-rainbow copy of P in $V(G) \setminus W^1$, so the minimum must be realized and the equalities must all be true. In particular, this means that $p_1 = p_3$, $p_2 = p_4$, and the edges $a_1^{p_1} a_2^{p_2}$ and $b_1^{p_3} b_2^{p_4} = b_1^{p_1} b_2^{p_2}$ are present in G, so $p := f^{-1}(p_1, p_2)$ satisfies $\Xi_p = \Psi_p = 1$, as claimed.

5.2 Comments on Connectedness

We conclude this section by briefly commenting on the requirement that P is connected in Theorem 5.2, as it seems that this is a noticeable drawback when compared to Theorem 5.1. We believe, however, that this is not as significant as it appears. Recall that these results are both for the CONGEST model. Now, if we allow that G and P are both disconnected, then the problem is impossible (as we could have one component of P per component of G and the nodes cannot communicate this to each other), so we must require that at least one of G and P is connected. By requiring that P is connected, we can allow G to be disconnected, since G contains a copy of P if and only if one of the connected components of G contains a copy of P. This aligns with Drucker et al. [11] in their proof for odd cycles, where they implicitly assume that their construction for G may be disconnected, as they add isolated vertices.

Chapter 6

Necklace Graphs

Given the results for rainbow subgraphs that we saw in the previous chapter, one might wonder if we could apply a similar strategy to show that, if P is a graph that contains an induced odd cycle of length at least 5, then detecting a copy of P takes at least $\Omega(n/\log n)$ rounds in the CONGEST model. We resolve this question in the special case of necklace graphs. Observe that, in the proofs for rainbow subgraphs, when we had a copy of P contained in the lower bound graph G, we were able to make use of a clever colouring scheme to significantly restrict where the copy could be. Without colours, we have to do significantly more work to achieve the same result, even just for necklace graphs, which have a very restricted structure. In this way, we provide motivation that a similar result might be true more generally in the uncoloured case (i.e., for any P that contains an induced cycle) but might require significantly more effort.

6.1 Lower Bound

The goal of this section is to prove the following lower bound on the round complexity of any distributed subgraph detection algorithm when the pattern graph is a necklace graph that contains an induced cycle of odd length at least 5.

Theorem 6.1. *Let P be a fixed necklace graph that contains an odd cycle of length at least 5 and consider the subgraph detection problem in the CONGEST model on all graphs G with n vertices. Then any algorithm that solves the problem takes at least $\Omega(n/\log n)$ rounds in the worst case.*

Before we begin, we comment that, for the sake of notational simplicity, we will occasionally abuse notation by interchangeably referring to a cycle by its vertex set.

Let P be given by Theorem 6.1 and suppose that P contains an odd cycle of length at least 5. We will prove the lower bound via a reduction from 2-party set disjointness. First, however, we will define some notation for P, and we suggest that the reader consult the example illustrated in Figure 6.1a. Let $C_1, C_2, \ldots, C_m$ be the cycles of P so that, for all $s \in [m-1]$, C_s shares exactly one vertex with C_{s+1}. For each $s \in [m]$, let l_s be the length of C_s. Let $k \in [m]$ so that $l_k \geq 5$ is odd and minimized, and k is maximized. Now, for each $s \in [k-1]$, write C_s as $w_{s,1}, w_{s,2}, \ldots, w_{s,l_s}$ so that w_{s,l_s} is the vertex shared with C_{s+1}. We also write that the shared vertex is w_{s+1,σ_s} when we refer to it as part of C_{s+1}. Similarly, for each $s \in [m] \setminus [k]$, write C_s as $w_{s,1}, w_{s,2}, \ldots, w_{s,l_s}$ so that w_{s,l_s} is the vertex shared with C_{s-1}. We also write the shared vertex as w_{s-1,σ_s} when we refer to it as part of C_{s-1}. To label the vertices in C_k, intuitively speaking, we split the cycle into two paths whose length differs by exactly 1 (since l_k is odd) so that the vertex shared by C_{k-1} and C_k is in one of the paths and the vertex shared by C_k and C_{k+1} is in the other (whenever these cycles exist). Note that we label the endpoints of these paths differently than the internal vertices, as this will be useful to us. Write the first path as $a_1, u_1, u_2, \ldots, u_{l_{k,1}}, b_1$ and the second path as $a_2, v_1, v_2, \ldots, v_{l_{k,2}}, b_2$, so that $a_1a_2 \in E(P)$, $b_1b_2 \in E(P)$, and $|l_{k,1} - l_{k,2}| = 1$. Now, let $u_0 = a_1$, $u_{l_{k,1}+1} = b_1$, $v_0 = a_2$, and $v_{l_{k,2}+1} = b_2$, as these additional labels will also be useful. Thus, we choose these paths with the additional condition that there is $\sigma_{k,1} \in [l_{k,1}+1] \cup \{0\}$ so that the vertex shared by C_{k-1} and C_k is $u_{\sigma_{k,1}}$ (when C_{k-1} exists) and that there is $\sigma_{k,2} \in [l_{k,2}+1] \cup \{0\}$ so that the vertex shared by C_k and C_{k+1} is $v_{\sigma_{k,2}}$ (when C_{k+1} exists). For simplicity, we choose $\sigma_{k,1} = 0$, which then determines the rest of the labels.

Now we will describe the reduction, and we suggest that the reader consult the example illustrated in Figure 6.1b throughout. Let N be an arbitrary positive integer and consider an instance of set disjointness of size N^2. Let Ξ and Ψ be the inputs to Alice and Bob, respectively. We will now describe how Alice and Bob construct a graph G that contains a copy of P if and only if there is $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$. As a general principle, given any vertex $z \in V(P)$, we will refer to various copies of z in G using integer superscripts, e.g., z^1, z^2, etc. We will describe the construction inductively. For any $i, j \in [m]$ so that $i \le k \le j$, let $P(i,j) = P[V(C_i) \cup \cdots \cup V(C_j)]$. For each such pair (i,j), we will describe how to construct the lower bound graph $G(i,j,\Xi,\Psi)$. The base case occurs when $i = j = k$. In this case, create N identical paths of length $l_{k,1}+2$, each corresponding to the vertices $a_1, u_1, u_2, \ldots, u_{l_{k,1}}, b_1$ from C_k and create N identical paths of length $l_{k,2}+2$, each corresponding to the vertices $a_2, v_1, v_2, \ldots, v_{l_{k,2}}, b_2$ from C_k. The endpoints of these paths will be joined by edges dictated by Alice's and Bob's inputs. More precisely, for all $p \in [N]$, add vertices a_1^p, b_1^p, a_2^p, and b_2^p to $G(i, j, \Xi, \Psi)$. These vertices are the endpoints of the paths, with the 'a' vertices being on Alice's "side" of the lower bound graph and the 'b' vertices being on Bob's "side" of the lower bound graph. Next, for all $p \in [N]$ and $q \in [l_{k,1}]$, add vertices u_q^p and for all $p \in [N]$ and $q \in [l_{k,2}]$, add vertices v_q^p. The 'u' and 'v' vertices will act as the internal nodes on the paths between the 'a' and 'b' vertices. Now we add edges. First, for all $p \in [N]$, connect the vertices $a_1^p, u_1^p, \ldots, u_{l_{k,1}}^p, b_1^p$ in a path. Similarly, for all $p \in [N]$, connect the vertices $a_2^p, v_1^p, v_2^p, \ldots, v_{l_{k,2}}^p, b_2^p$ in a path. Finally, we describe how Alice and Bob add the edges associated with their input bit strings. Notice that, excluding the edges added in this step, the entire construction is known to both Alice and Bob. First, fix an arbitrary bijection $f : [N^2] \to [N] \times [N]$. To slightly abuse notation, we will write that, for each $p \in [N^2]$, $f(p) = (p_1, p_2)$. Then, for each $p \in [N^2]$ so that $\Xi_p = 1$, Alice computes $f(p)$ to obtain (p_1, p_2) and then adds an edge from $a_1^{p_1}$ to $a_2^{p_2}$ in G. Similarly, for each $p \in [N^2]$ so that $\Psi_p = 1$, Bob computes $f(p)$ to obtain (p_1, p_2) and then adds an edge from $b_1^{p_1}$ to $b_2^{p_2}$ to G. Now, if the necklace consisted of exactly one cycle, then the construction would be complete, and observe that the constructed graph contains C_k if and only if there is an $p \in [N^2]$ such that $\Xi_p = \Psi_p = 1$, as illustrated in Chapter 4.

Next, we consider the inductive case. Let $i, j \in [m]$ be so that $i \leq k \leq j$ and assume that we have already defined the graph $G(i, j, \Xi, \Psi)$. There are two lower bound graphs that we define: (1) $G(i-1, j, \Xi, \Psi)$ (when $i > 1$) and (2) $G(i, j+1, \Xi, \Psi)$ (when $j < m$). Intuitively, $G(i-1, j, \Xi, \Psi)$ is obtained by "attaching" C_{i-1} to $G(i, j, \Xi, \Psi)$ and $G(i, j+1, \Xi, \Psi)$ is obtained by "attaching" C_{j+1} to $G(i, j, \Xi, \Psi)$.

The precise definitions of these graphs are as follows, described separately.

(1) Begin by adding vertices $w^1_{i-1,1}, w^1_{i-1,2}, \ldots, w^1_{i-1,l_{i-1}-1}$ to $G(i,j,\Xi,\Psi)$ and connect them as a path. There are two possibilities based on the value of i. First, consider the case where $i = k$. Then, for each $p \in [N]$, add an edge between $w^1_{i-1,1}$ and a^p_1 and between $w^1_{i-1,l_{i-1}-1}$ and a^p_1. In the second case, consider $i < k$. Then, add an edge between $w^1_{i-1,1}$ and $w^1_{i,\sigma_{i-1}}$ and between $w^1_{i-1,l_{i-1}-1}$ and $w^1_{i,\sigma_{i-1}}$.

(2) Begin by adding vertices $w^1_{j+1,1}, w^1_{j+1,2}, \ldots, w^1_{j+1,l_{j+1}-1}$ to $G(i,j,\Xi,\Psi)$ and connect them as a path. There are two possibilities based on the value of j. First, consider the case where $j = k$. Then, for each $p \in [N]$, add an edge between $w^1_{j+1,1}$ and $v^p_{\sigma_{k,2}}$ and between $w^1_{j+1,l_{j+1}-1}$ and $v^p_{\sigma_{k,2}}$. In the second case, consider $j > k$. Then, add an edge between $w^1_{j+1,1}$ and $w^1_{j,\sigma_{j+1}}$ and between $w^1_{j+1,l_{j+1}-1}$ and $w^1_{j,\sigma_{j+1}}$.

Thus for any pair $i,j \in [N]$ so that $i \leq k \leq j$, we can construct the graph $G(i,j,\Xi,\Psi)$ by starting with $G(k,k,\Xi,\Psi)$ and appropriately applying the two inductive constructions described above. We conclude the construction by defining $G = G(1,m,\Xi,\Psi)$.

In the following proofs, we will make use of the following notation related to the lower bound graph G. For all $s \in [m] \setminus \{k\}$, let $W_s = \{w^1_{s,q} : q \in [l_s - 1]\}$. Each set W_s represents a copy of the vertices of the cycle C_s from the necklace graph P. Note that, intuitively speaking, we have excluded the last vertex of the cycle since it is shared with the "next" cycle "closer" to C_k. Let $A = \{a^p_q : q \in [2], p \in [N]\}$ and $B = \{b^p_q : q \in [2], p \in [N]\}$. Let $M = V(G) \setminus \bigcup_{s \in [m]\setminus\{k\}} W_s$. The set M consists of the vertices along the $2N$ parallel paths between the A and B vertices, i.e., the paths that were added in the construction of $G(k,k,\Xi,\Psi)$.

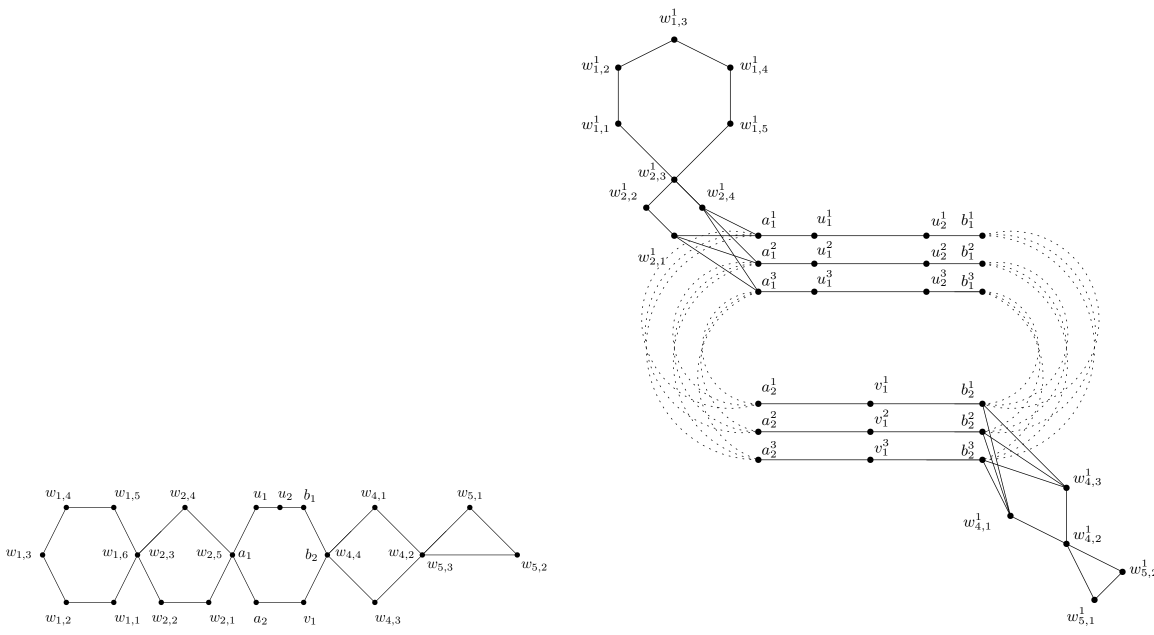

(a) (b)

Figure 6.1: (a) a necklace graph P. (b) a lower bound graph G constructed for P with $N = 3$, $i = 1$, $j = 5$, and $k = 3$. The dotted edges represent all the possible edges that could be added depending on the values of Ξ and Ψ.

The core of the proof of Theorem 6.1 is Lemma 6.2. At a high level, Lemma 6.2 says that if G contains a copy of P, then there is a copy of C_k in $G[M]$, i.e., in the subgraph of G induced by the vertices of M.

Lemma 6.2. *If $G(1, m, \Xi, \Psi)$ contains a copy of $P(1, m)$, then there is a copy of $P(k, k)$ in $G(k, k, \Xi, \Psi)$.*

The proof of Lemma 6.2 can be found in Section 6.1.2. Using Lemma 6.2, we now prove Theorem 6.1.

Proof of Theorem 6.1. We start by showing that that G contains a copy of P if and only if there is an $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$. First, suppose there is an $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$, and suppose $f(p) = (p_1, p_2)$. Then Alice will add an edge between $a_1^{p_1}$ and $a_2^{p_2}$ and Bob will add an edge between $b_1^{p_1}$ and $b_2^{p_2}$, so the subgraph of G induced by $\{a_1^{p_1}, u_1^{p_1}, \ldots, u_{l_{k,1}}^{p_1}, b_1^{p_1}, b_2^{p_2}, v_{l_{k,2}}^{p_2}, \ldots, v_1^{p_2}, a_2^{p_2}\}$ is a copy of C_k. Then, adding the vertices from $\bigcup_{s\in[m]\setminus\{k\}} W_s$ gives a copy of P.

On the other hand, suppose that G contains a copy of P. We will show that there is a $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$. By Lemma 6.2, there is a copy of C_k in $G[M]$. By design, the graph $G[M]$ is exactly the lower bound graph constructed by Drucker et al. [11]. Thus, by Lemma 4.1, we find that there is a $p \in [N^2]$ so that, if $f(p) = (p_1, p_2)$, then the edges $a_1^{p_1}a_2^{p_2}$ and $b_1^{p_1}b_2^{p_2}$ are in G. Thus, by the construction of G, $\Xi_p = \Psi_p = 1$, as claimed.

Now, we determine which nodes Alice and Bob will simulate. Intuitively, we choose the cut edges to be the rightmost edges of the paths from A vertices to B vertices, so everything to the "left" of the B nodes will be simulated by Alice and everything at or to the "right" of the B nodes will be simulated by Bob. So, in most cases, the B nodes will be simulated by Bob and everything else will be simulated by Alice, however, in the case where the copies of $C_{k+1}, \ldots, C_m$ are attached via the B nodes (as in Figure 6.1b), then the nodes in these copies will be simulated by Bob as well. Now, we make this precise. Let $L = \bigcup_{s=1}^{k-1} W_s$ and $R = \bigcup_{s=k+1}^{m} W_s$. Then there are two possibilities based on the vertex shared by C_k and C_{k+1}.

- If the vertex shared by C_k and C_{k+1} is b_2, then Bob simulates the nodes in $B \cup R$ and Alice simulates the nodes in $(M \setminus B) \cup L$.
- Otherwise, Bob simulates the nodes in B and Alice simulates the nodes in $(M \setminus B) \cup L \cup R$.

Finally, it remains to compute the number of nodes in the network and the size of the cut between Alice and Bob. First, let n be the number of vertices of G. Notice that we have made N copies of each of the nodes in C_k, giving $N \cdot l_k$ vertices. Then we include one copy of each node of P except those in C_k, which gives $|V(P)| - l_k$ vertices. Thus, we have that $n = N \cdot l_k + |V(P)| - l_k = (N-1) \cdot l_k + |V(P)|$. Note that this means that $n \in \Theta(N)$, as N is the only non-constant value in this expression, so we have that $N \in \Theta(n)$. Now notice that there are precisely $2N$ cut edges, as they are exactly the edges $u^p_{l_{k,1}} b^p_1$ and $v^p_{l_{k,2}} b^p_2$ for all $p \in [N]$. Therefore, as described in Chapter 4, we derive a lower bound of

$$\frac{N^2 + 1 - 1}{2 \cdot 2N \log n} = \frac{N^2}{4N \log n} = \frac{N}{4 \log n} \in \Omega\left(\frac{n}{\log n}\right)$$

as claimed.

6.1.1 Spotting a Necklace in G

In this section, we characterize what a copy of P in G can look like. We will use the following three results in our proof of Lemma 6.2 in Section 6.1.2. Our first result, Lemma 6.3, confirms that, outside of $G[W_{k-1} \cup M \cup W_{k+1}]$, the lower bound graph G replicates the necklace structure of P.

Lemma 6.3. *Let $s \in [k-2]$ and let C be a cycle of G that contains a vertex from W_s but no vertex from W_{s-1}. Then $V(C) = W_s \cup \{w^1_{s+1,\sigma_s}\}$. Similarly, let $s \in [m] \setminus [k+1]$ and let C be a cycle of G that contains a vertex from W_s but no vertex from W_{s+1}. Then $V(C) = W_s \cup \{w^1_{s-1,\sigma_s}\}$.*

Proof. Without loss of generality, we assume that $s \in [k-2]$, as the proof when

$s \in [m]\backslash[k+1]$ is symmetric. Let $z_1 \in V(C) \cap W_s$. Recall that $W_s = \{w^1_{s,1}, \ldots, w^1_{s,l_s-1}\}$, and, from the inductive case in the definition of G, that the vertices of W_s are connected as a path with endpoints $w^1_{s,1}$ and w^1_{s,l_s-1}. So, since $z_1 \in W_i$, we know that z_1 is some vertex along with path with endpoints $w^1_{s,1}$ and w^1_{s,l_s-1}. Next, since C is a cycle, C consists of two internally disjoint paths that start at z_1 and share a common endpoint, which we denote by z_2. Since $z_1 \in V(C)$ and by the assumption that C contains no vertices from W_{s-1}, the only two disjoint paths with common starting vertex z_1 are the paths from z_1 to $w^1_{s,1}$ and from z_1 to w^1_{s,l_s-1} in W_s. Note that the union of these two paths is exactly the path formed by the nodes of W_s. Next, from the the inductive case in the definition of G, note that the only vertex adjacent to $w^1_{s,1}$ outside of W_s is w^1_{s+1,σ_s} and the only vertex adjacent to w^1_{s,l_s-1} outside of W_s is w^1_{s+1,σ_s}. Thus, $z_2 = w^1_{s+1,\sigma_s}$, which completes the proof that $V(C) = W_s \cup \{w^1_{s+1,\sigma_s}\}$, as claimed. □

Now we consider what a copy of $P(i,j)$ in $G(i,j,\Xi,\Psi)$ can look like. Roughly speaking, in Lemma 6.4 we observe that the leftmost and rightmost cycles of $G(i,j,\Xi,\Psi)$ can only be mapped onto by the leftmost and rightmost cycles of $P(i,j)$, assuming that such cycles of $G(i,j,\Xi,\Psi)$ are not too close to $G[M]$, i.e., there is at least one cycle "between" them and $G[M]$.

Lemma 6.4. *Let $i,j \in [m]$ be so that $i \leq k \leq j$ and $j - i \geq 2$ and suppose that $s \in \{i,j\}$ such that $|s-k| \geq 2$. Suppose further that $G(i,j,\Xi,\Psi)$ contains a copy of $P(i,j)$ via the subgraph mapping $\phi : V(P(i,j)) \to V(G(i,j,\Xi,\Psi))$. Then exactly one of the following is true:*

1. $\phi^{-1}(W_s) = \emptyset$

2. $\emptyset \neq \phi^{-1}(W_s) \subseteq \{w_{i,1}, w_{i,2}, \ldots, w_{i,l_i-1}\}$

3. $\emptyset \neq \phi^{-1}(W_s) \subseteq \{w_{j,1}, w_{j,2}, \ldots, w_{j,l_j-1}\}$

Proof. Without loss of generality, we assume that $s = i$, as the proof when $s = j$ is symmetric. It suffices to show that the following two situations are impossible: (1) there is $t \in \{i+1, \ldots, j-1\}$ so that $\phi(V(C_t))$ contains a vertex from W_i and (2) both $\phi(V(C_i))$ and $\phi(V(C_j))$ contain a vertex from W_i. To see that this suffices, notice that, by showing that situation (1) is impossible, we show that only vertices from $V(P(i,j)) \setminus V(C_{i+1}) \cup \cdots \cup V(C_{j-1}) = \{w_{i,1}, w_{i,2}, \ldots, w_{i,l_i-1}, w_{j,1}, w_{j,2}, \ldots, w_{j,l_j-1}\}$ can map to vertices in W_i. Next, showing that situation (2) is impossible shows that vertices from at most one of the sets $\{w_{i,1}, w_{i,2}, \ldots, w_{i,l_i-1}\}$ or $\{w_{j,1}, w_{j,2}, \ldots, w_{j,l_j-1}\}$ can map to vertices in W_i. Finally, combining these gives that vertices from at most one of the sets $\{w_{i,1}, w_{i,2}, \ldots, w_{i,l_i-1}\}$ or $\{w_{j,1}, w_{j,2}, \ldots, w_{j,l_j-1}\}$ can map to vertices in W_i and that no other vertices from $V(P(i,j))$ can map to vertices in W_i, which is exactly what we need to show.

(1) Assume that there is a $t \in \{i+1, \ldots, j-1\}$ so that $\phi(V(C_t))$ contains a vertex from W_i. Since $\phi(V(C_t))$ is a subset of $V(G(i,j,\Xi,\Psi))$, which consists of the vertices $M \cup W_i \cup \cdots \cup W_j$, we may apply Lemma 6.3 to conclude that $\phi(V(C_t)) = W_i \cup \{w^1_{i+1,\sigma_i}\}$. From the choice of t, we know that C_t is a cycle in $P(i,j)$ that is not C_i and not C_j, so C_t shares a vertex with exactly two cycles in $P(i,j)$. It follows that $\phi(V(C_t))$ must share a vertex with at least two cycles in $G(i,j,\Xi,\Psi)$. However, in $G(i,j,\Xi,\Psi)$, $W_i \cup \{w^1_{i+1,\sigma_i}\}$ shares a vertex with exactly one other cycle, namely, it shares the vertex w^1_{i+1,σ_i} with the cycle that contains all of W_{i+1}. Thus, we obtain a contradiction, so situation (1) is impossible.

(2) Assume that $\phi(V(C_i))$ and $\phi(V(C_j))$ both contain a vertex from W_i. Since $\phi(V(C_i))$ and $\phi(V(C_j))$ are subsets of $V(G(i,j,\Xi,\Psi))$, which consists of the vertices $M \cup W_i \cup \cdots \cup W_j$, we may apply Lemma 6.3 to conclude that $\phi(V(C_i)) = W_i \cup \{w^1_{i+1,\sigma_i}\}$ and $\phi(V(C_j)) = W_i \cup \{w^1_{i+1,\sigma_i}\}$. Thus, $\phi(V(C_i)) = \phi(V(C_j))$. However, since $j - i \geq 2$, C_i and C_j do not share any vertices, and hence $\phi(V(C_i)) \cap \phi(V(C_j)) = \emptyset$, which contradicts $\phi(V(C_i)) = \phi(V(C_j))$. Thus, situation (2) is impossible. $\square$

Finally, in Lemma 6.5 we observe that, if one of the ends of $G(i,j,\Xi,\Psi)$ is mapped onto by the opposite end of $P(i,j)$ (for example, the cycle C_i in $P(i,j)$ gets mapped to the copy of the cycle C_j in $G(i,j,\Xi,\Psi)$), then in the mapping of $P(i,j)$ to $G(i,j,\Xi,\Psi)$, the cycles in the necklace will appear in the reverse order up to whichever of W_{k-2} or W_{k+2} is reached first in $G(i,j,\Xi,\Psi)$. See Figure 6.2 for an illustration of this.

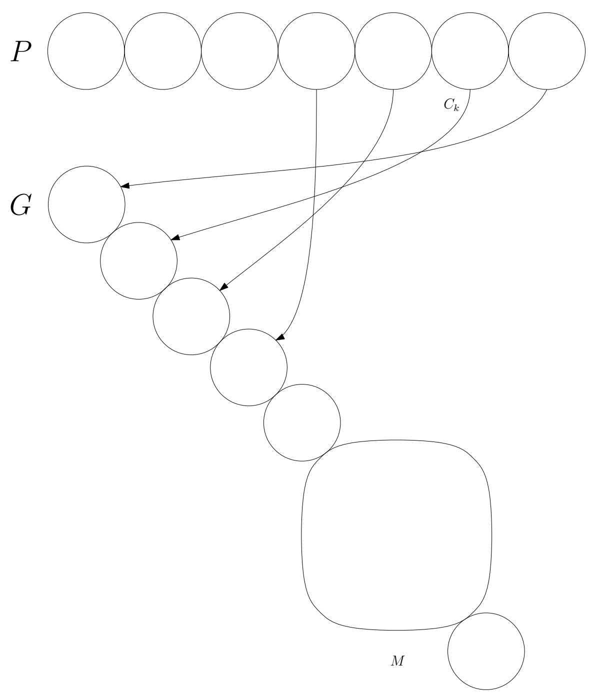

Figure 6.2: An illustration of Lemma 6.5. The black arrows represents the mapping ϕ.

Lemma 6.5. *Let $i, j \in [m]$ be so that $i \leq k \leq j$. Suppose that $G(i,j,\Xi,\Psi)$ contains a copy of $P(i,j)$ via the subgraph mapping $\phi : V(P(i,j)) \to V(G(i,j,\Xi,\Psi))$. Then the following are true:*

(1) Suppose that $\phi^{-1}(W_i) \subseteq \{w_{j,1}, w_{j,2}, \ldots, w_{j,l_j-1}\}$. Then, for all $s \in \{i, \ldots, k-2\}$, $\phi(V(C_{j-(s-i)})) = W_s \cup \{w^1_{s+1,\sigma_s}\}$.

(2) Suppose that $\phi^{-1}(W_j) \subseteq \{w_{i,1}, w_{i_2}, \ldots, w_{i,l_i-1}\}$. Then, for all $s \in \{k+2, \ldots, j\}$, $\phi(V(C_{i+(j-s)}) = W_s \cup \{w^1_{s-1,\sigma_s}\}$.

Proof. We prove statement (1), as the proof of statement (2) is symmetric.

The proof is by induction on s.

Base case: $s = i$. Since $\phi(V(C_j))$ is a cycle in $G(i,j,\Xi,\Psi)$ and we assume that $\phi^{-1}(W_i) \subset V(C_j)$, it follows that $\phi(V(C_j))$ contains a vertex from W_i. From the inductive construction of $G(i,j,\Xi,\Psi)$, the vertices of $G(i,j,\Xi,\Psi)$ consist of $M \cup W_i \cup \cdots \cup W_j$, that is, $\phi(V(C_j))$ does not contain any vertices of W_{i-1}. Thus, we apply Lemma 6.3 to get that $\phi(V(C_j)) = W_i \cup \{w^1_{i+1,\sigma_i}\}$, as claimed.

Induction hypothesis: Let $s \in \{i+1, \ldots, k-2\}$ and suppose that, for all $s' \in \{i, \ldots, s-1\}$, $\phi(V(C_{j-(s'-i)})) = W_{s'} \cup \{w^1_{s'+1,\sigma_{s'}}\}$.

Inductive step: We set out to prove that $\phi(V(C_{j-s+i})) = W_s \cup \{w^1_{s+1,\sigma_s}\}$. By the indexing of the cycles in the necklace $P(i,j)$, we have that $C_{j-(p-i)}$ shares exactly one vertex with $C_{j-((p-1)-i)}$ and no vertices with $C_{j-(p''-i)}$ for each $p'' \in \{i, \ldots, p-2\}$. Thus, under the mapping ϕ, the cycles must obey the same structure, i.e., it follows that $\phi(V(C_{j-(s-i)}))$ forms a cycle in $G(i,j,\Xi,\Psi)$ that shares exactly one vertex with $\phi(V(C_{j-((s-1)-i)}))$ and no vertices with $\phi(V(C_{j-(s''-i)}))$, for any $s'' \in \{i, \ldots, s-2\}$. By applying the induction hypothesis for each $s' \in \{i, \ldots, s-1\}$, this means that $\phi(V(C_{j-(s-i)}))$ shares exactly one vertex with $W_{s-1} \cup \{w^1_{s,\sigma_{s-1}}\}$ (which means that it contains at most one vertex from W_{s-1}) and no vertices with any of $W_i, W_{i+1}, \ldots, W_{s-2}$. Finally, we show that $\phi(V(C_{j-(s-i)})$ does not contain any vertices from W_{s-1}. Suppose, for the sake of a contradiction, that $\phi(V(C_{j-(s-i)})$ contains a vertex $w^1_{s-1,q} \in W_{s-1}$. Now, in $G(i, j, \Xi, \Psi)$, $w^1_{s-1,q}$ only has neighbours in W_{s-2}, W_{s-1}, and W_s. However, since $\phi(V(C_{j-(s-i)})$ contains at most one vertex from W_{s-1} and no vertices from W_{s-2}, both neighbours of $w^1_{s-1,q}$ in $\phi(V(C_{j-(s-i)})$ are in W_s. However, this is impossible since the fact that $s \neq k$ gives that each vertex of W_{s-1} is adjacent to at most one vertex in W_s (namely, the vertex $w^1_{s,\sigma_{s-1}}$)). Thus, since we have found a contradiction, we have that $\phi(V(C_{j-s+i}))$ forms a cycle in $G(i, j, \Xi, \Psi)$ that contains a vertex from W_s and no vertex from W_{s-1}, so we apply Lemma 6.3 to get that $\phi(V(C_{j-s+i})) = W_s \cup \{w^1_{s+1,\sigma_s}\}$, as claimed.

Note that, in the analogous proof of statement (2), the base case happens when $s = j$ and the induction hypothesis assumes that the statement is true for all $s' \in \{s+1, \ldots, j\}$. $\square$

6.1.2 Proof of Lemma 6.2

Lemma 6.2. *If $G(1, m, \Xi, \Psi)$ contains a copy of $P(1, m)$, then there is a copy of $P(k, k)$ in $G(k, k, \Xi, \Psi)$.*

To prove Lemma 6.2, we will prove the following statement: for any $i, j \in [m]$ such that $i \leq k \leq j$, if $G(i, j, \Xi, \Psi)$ contains a copy of $P(i, j)$, then $G(k, k, \Xi, \Psi)$ contains a copy of $P(k, k)$. Setting $i = 1$ and $j = m$ finishes the proof. With regards to the structure of the proof, the current subsection contains the significant steps, and the minor technical results have been relegated to Section 6.1.4.

We prove Lemma 6.2 by induction on the values of i and j. The first base case occurs when $i = j = k$, in which case the statement is trivially true. The second base case occurs when $j - i = 1$, and its proof is in Section 6.1.3.1. The third base case occurs when $j - i = 2$ and $k = i + 1$, and its proof is in Section 6.1.3.2. As induction hypothesis, suppose that $j - i \geq 2$ and assume that, for any $i', j' \in [m]$ such that $i' \leq k \leq j'$ and $j' - i' < j - i$, if $G(i', j', \Xi, \Psi)$ contains a copy of $P(i', j')$ then $G(k, k, \Xi, \Psi)$ contains a copy of $P(k, k)$. To prove the inductive step, we suppose that $G(i, j, \Xi, \Psi)$ contains a copy of $P(i, j)$, and we set out to prove that $G(k, k, \Xi, \Psi)$ contains a copy of $P(k, k)$. Note that, from our choice of base cases, we may assume that $j - i \geq 2$ and there exists an $s \in \{i, j\}$ such that $|s - k| \geq 2$. For any such s, we consider 3 cases: (1) $\phi^{-1}(W_s) = \emptyset$; (2) $\emptyset \neq \phi^{-1}(W_s) \subseteq \{w_{s,1}, \ldots, w_{s,l_s-1}\}$; and (3) $\emptyset \neq \phi^{-1}(W_s) \subseteq \{w_{t,1}, \ldots, w_{t,l_t-1}\}$ where $t = \{i, j\} \setminus \{s\}$. These 3 cases are exhaustive, by Lemma 6.4. We now prove that $G(k, k, \Xi, \Psi)$ contains a copy of $P(k, k)$ in each of the three cases:

Case 1: Suppose that $\phi^{-1}(W_s) = \emptyset$. We focus on the case when $s = i$, as the proof when $s = j$ is symmetric. Due to the assumption that $\phi^{-1}(W_i) = \emptyset$, it follows that $\phi(V(P(i, j)) \subseteq V(G(i + 1, j, \Xi, \Psi))$, so $\phi(V(P(i + 1, j)) \subseteq V(G(i + 1, j, \Xi, \Psi))$. From the fact that $|s - k| \geq 2$, it follows that $i + 1 \leq k \leq j$, so we can apply the induction hypothesis with $i' = i + 1$ and $j' = j$ to get that $G(k, k, \Xi, \Psi)$ contains a copy of $P(k, k)$, as claimed.

Case 2: Suppose that $\emptyset \neq \phi^{-1}(W_s) \subseteq \{w_{s,1}, w_{s,2}, \ldots, w_{s,l_s-1}\}$. We focus on the case when $s = i$, as the proof when $s = j$ is symmetric. From the definitions of $P(i,j)$ and $G(i,j,\Xi,\Psi)$, we have that, since $V(P(i,j)) \setminus \{w_{i,1}, \ldots, w_{i,l_i-1}\} = V(P(i+1,j))$ and $V(G(i,j,\Xi,\Psi)) \setminus W_i = V(G(i+1,j,\Xi,\Psi))$, so from the fact that $\phi^{-1}(W_s) \subseteq \{w_{s,1}, \ldots, w_{s,l_s-1}\}$, it follows that $\phi(V(P(i+1,j))) \subseteq V(G(i+1,j,\Xi,\Psi))$. From the fact that $|s-k| \geq 2$, it follows that $i+1 \leq k \leq j$, so we can apply the induction hypothesis with $i' = i+1$ and $j' = j$ to get that $G(k,k,\Xi,\Psi)$ contains a copy of $P(k,k)$, as claimed.

Case 3: Let $t = \{i,j\} \setminus \{s\}$ and suppose that $\emptyset \neq \phi^{-1}(W_s) \subseteq \{w_{t,1}, w_{t,2}, \ldots, w_{t,l_t-1}\}$. Since $i \leq k \leq j$, exactly one of the following must be true: $k \in \{i+2, i+3, \ldots, j-3, j-2\}$, $k \in \{i,j\}$, or $k \in \{i+1, j-1\}$. We consider the first two of these subcases below as Case 3.1 and Case 3.2, respectively. We separate the case where $k \in \{i+1, j-1\}$ into two further subcases that occur when $j - i \geq 4$ (considered in Case 3.3 below) and when $j - i = 3$ (considered in Case 3.4 below). Note that we do not need to consider $j - i \leq 2$ as such a situation is handled by the three base cases.

Case 3.1: Suppose that $k \in \{i+2, \ldots, j-2\}$. Then, by Lemma 6.4, $\phi^{-1}(W_i \cup W_j) \subseteq \{w_{i,1}, \ldots, w_{i,l_i-1}, w_{j,1}, \ldots, w_{j,l_j-1}\}$. From the definitions of $P(i,j)$ and $G(i,j,\Xi,\Psi)$, we know that $V(P(i,j)) \setminus \{w_{i,1}, \ldots, w_{i,l_i-1}, w_{j,1}, \ldots, w_{j,l_j-1}\} = V(P(i+1,j-1))$ and $V(G(i,j,\Xi,\Psi)) \setminus (W_i \cup W_j) = V(G(i+1,j-1,\Xi,\Psi))$. Since $k \in \{i+2, \ldots, j-2\}$, we can apply the induction hypothesis with $i' = i+1$ and $j' = j-1$ to get that $G(k,k,\Xi,\Psi)$ contains a copy of $P(k,k)$, as claimed.

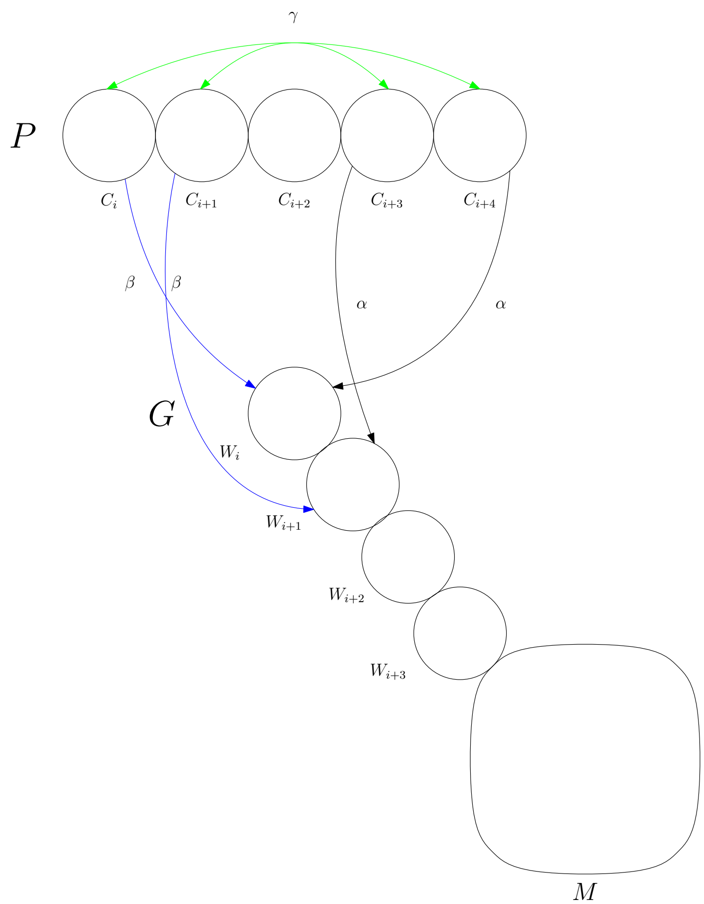

Figure 6.3: Illustration of Case 3.2. The black arrows (those labelled α) represent cycles that we know have the same length by ϕ and Lemma 6.5. The blue arrows (those labelled β) represent cycles that we know have the same length by construction. The green arrows (those labelled γ) represent cycles that we can thus conclude have the same length.

Case 3.2: Suppose that $k \in \{i, j\}$. We focus on the case when $k = j$, as the proof when $k = i$ is symmetric. Recall that $j - i \geq 2$ by assumption, so $j - i + 1 \geq 3$. Thus, by Lemma 6.29 (found in Section 6.1.4.5), $\lfloor \frac{j-i+1}{2} \rfloor \leq j - i - 1$, which means that $\lfloor \frac{j-i+1}{2} \rfloor - 1 \leq j - i - 2$. Now, for all $s \in \{i, \ldots, \lfloor \frac{i+j-1}{2} \rfloor\}$, let $d \in \left[\lfloor \frac{j-i+1}{2} \rfloor - 1\right] \cup \{0\}$ be so that $s = i + d$. Then $s = i + d \leq i + \left(\lfloor \frac{j-i+1}{2} \rfloor - 1\right) \leq i + (j - i - 2) = j - 2 = k - 2$. Thus, by Lemma 6.5, we have that $\phi(V(C_{j-s+i})) = W_s \cup \{w^1_{s+1,\sigma_s}\}$, so $l_s = |W_s \cup \{w^1_{s+1,\sigma_1}\}| = |\phi(V(C_{j-s+i}))| = l_{j-s+i}$. However, since this is true for all $s \in \{i, \ldots, \lfloor \frac{i+j-1}{2} \rfloor\}$, this means that $P(i+1, j)$ is isomorphic to $P(i, j-1)$. Intuitively, this is because we have proved that the graph $P(i, j)$ is symmetric. See Figure 6.3 for an illustration of this.

Now, since $G(i, j, \Xi, \Psi)$ contains a copy of $P(i, j)$ via ϕ and we know that $\phi(V(C_j)) = W_i \cup \{w^1_{i+1,q_i}\}$ (applying Lemma 6.5 when $s = i$), it follows that $G(i+1, j, \Xi, \Psi)$ contains a copy of $P(i, j-1)$. However, since $P(i+1, j)$ is isomorphic to $P(i, j-1)$, this means that $G(i+1, j, \Xi, \Psi)$ contains a copy of $P(i+1, j)$ as well. Since $j - i \geq 2$ by assumption, we have that $i + 1 \leq k = j$, so we can apply the induction hypothesis with $i' = i + 1$ and $j' = j$ to conclude that $G(k, k, \Xi, \Psi)$ contains a copy of $P(k, k)$, as claimed.

Case 3.3: Suppose that $j - i \geq 4$ and $k \in \{i+1, j-1\}$. We will focus on the case when $k = j-1$, as the proof when $k = i+1$ is symmetric. Since $j-i \geq 4$, we have that $j-i+1 \geq 5$, so Lemma 6.30 (found in Section 6.1.4.5) gives that $\lfloor \frac{j-i+1}{2} \rfloor \leq j-i-2$, so $\lfloor \frac{j-i+1}{2} \rfloor - 1 \leq j - i - 3$. Now, for all $s \in \{i, \ldots, \lfloor \frac{i+j-1}{2} \rfloor\}$, let $d \in \left[\lfloor \frac{j-i+1}{2} \rfloor - 1\right] \cup \{0\}$ be so that $s = i + d$. Then $s = i + d \leq i + \left(\lfloor \frac{j-i+1}{2} \rfloor - 1\right) \leq i + (j - i - 3) = j - 3 = k - 2$. Thus, by Lemma 6.5, we have that $\phi(V(C_{j-s+i})) = W_s \cup \{w^1_{s+1,\sigma_s}\}$, so $l_s = |W_s \cup \{w^1_{s+1,\sigma_s}\}| = |\phi(V(C_{j-s+i}))| = l_{j-s+i}$. However, since this is true for all $s \in \{i, \ldots, \lfloor \frac{i+j-1}{2} \rfloor\}$, this means that $P(i+1, j)$ is isomorphic to $P(i, j-1)$ (as explained in Case 3.2).

Now, since $G(i, j, \Xi, \Psi)$ contains a copy of $P(i, j)$ via ϕ, and we know that $\phi(V(C_j)) = W_i \cup \{w^1_{i+1,q_i}\}$ (applying Lemma 6.5 when $s = i$), it follows that $G(i+1, j, \Xi, \Psi)$ contains a copy of $P(i, j-1)$. However, since $P(i+1, j)$ is isomorphic to $P(i, j-1)$, this means that $G(i+1, j, \Xi, \Psi)$ contains a copy of $P(i+1, j)$ as well. Since $j - i \geq 4$ by assumption, we get that $i + 1 \leq k = j - 1 \leq j$, so we can apply the induction hypothesis with $i' = i + 1$ and $j' = j$ to conclude that $G(k, k, \Xi, \Psi)$ contains a copy of $P(k, k)$, as claimed.

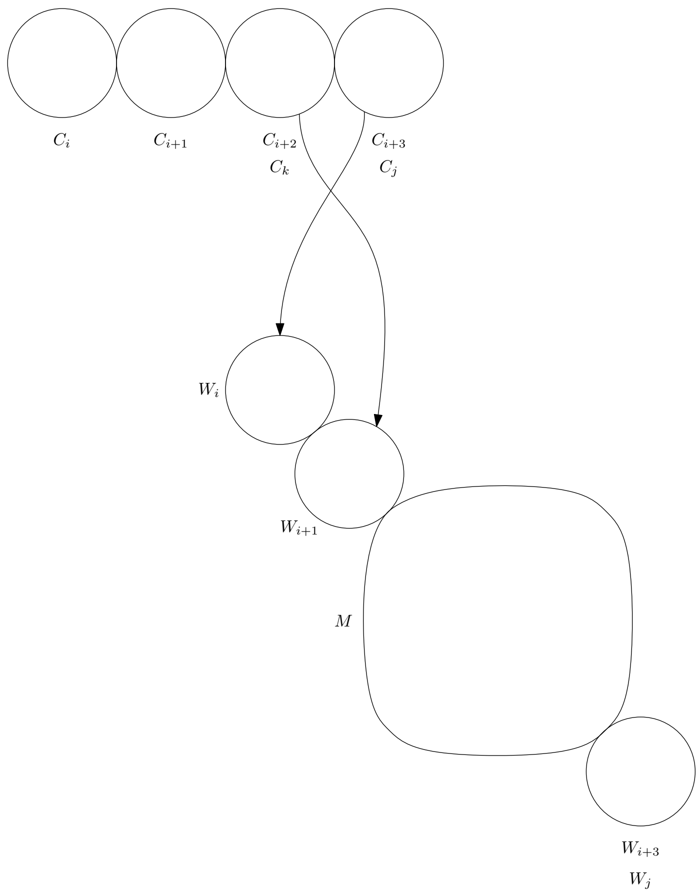

Figure 6.4: An illustration of Case 3.4. The arrows represent the mapping ϕ. Note that $\phi(V(C_{i+2}))$ need not contain all of W_{i+1}.

Case 3.4: Suppose that $j - i = 3$ and $k \in \{i+1, j-1\}$. We focus on the case when $k = j - 1$, as the proof when $k = i + 1$ is symmetric. First, by Lemma 6.5 with $s = i$, we have that $\phi(V(C_j)) = W_i \cup \{w^1_{i+1,\sigma_i}\}$. Thus, we see that $l_i = |W_i \cup \{w^1_{i+1,\sigma_i}\}| = |\phi(V(C_j))| = l_j$. Our goal is to prove that $l_{i+1} = l_{j-1}$ so that we can apply a similar argument as we did in Cases 3.2 and 3.3 (i.e., using the fact that the necklace $P(i,j)$ is symmetric). In those cases, we were able to use Lemmas 6.29 or 6.30 together with Lemma 6.5 to prove symmetry, and thus conclude that $P(i+1,j)$ is isomorphic to $P(i,j-1)$. However, since we are currently assuming that $j-i=3$, we have that $j-i+1=4$, so Lemmas 6.29 and 6.30 do not apply. We will require more careful arguments.

In the remainder of the proof, we use the fact that $j-1 = i+2$ (since we are assuming that $j-i=3$), so our goal of proving the symmetry of $P(i,j) = P(i,i+3)$ boils down to proving that the cycles C_{i+1} and C_{i+2} are of the same size, i.e., that $l_{i+1} = l_{i+2}$. Since $\phi(V(C_{i+3})) = W_i \cup \{w^1_{i+1,\sigma_i}\}$ and since $\phi(V(C_{i+3}))$ shares exactly one vertex with $\phi(V(C_{i+2}))$, we get that the shared vertex is w^1_{i+1,σ_i}. Thus, $\phi(V(C_{i+2}))$ contains at least one vertex from W_{i+1} and no vertices from W_i. Then, since $G[W_{i+1}]$ is a path and $\phi(V(C_{i+2}))$ is a cycle, it follows that $\phi(V(C_{i+2}))$ must contain at least one node not contained in W_{i+1}. Other than the nodes in W_i, the only nodes that are adjacent to nodes in W_{i+1} belong to M (by construction of $G(i,j,\Xi,\Psi)$), so it follows that $\phi(V(C_{i+2}))$ contains at least one node from M. An illustration of this can be found in Figure 6.4. Now, it is either the case that $\phi(V(C_{i+2}))$ contains no vertices from W_{i+3} or at least one vertex from W_{i+3}. We consider these cases separately:

Case 3.4.1: Suppose that $\phi(V(C_{i+2}))$ contains no vertices from W_{i+3}. Then Lemma 6.20 (in Section 6.1.4.3) gives that $\phi(V(C_{i+2}))$ contains all of W_{i+1}. We argued above that $\phi(V(C_{i+2}))$ contains at least one vertex from M, but we now show that it contains exactly one vertex from M. To see this, assume, for the sake of contradiction, that $\phi(V(C_{i+2}))$ contains at least two vertices from M. Note that this means that $l_{i+2} = |\phi(V(C_{i+2}))| \geq |W_{i+1}| + 2 = l_{i+1} + 1 > l_{i+1}$ and, by Lemma 6.20, l_{i+1} is odd. Thus, since $k = i + 2$ was defined such that l_k is the smallest odd cycle length that is at least 5, it follows that $l_{i+1} = 3$. Thus, $\phi(V(C_{i+1}))$ is a 3-cycle in $G(i, i + 3, \Xi, \Psi)$. Since $V(C_{i+1})$ shares exactly one vertex with $V(C_{i+2})$ and no vertices with $V(C_{i+3})$ in $P(i, i + 3)$, it follows that $\phi(V(C_{i+1}))$ shares exactly one vertex with $\phi(V(C_{i+2}))$ and no vertices with $\phi(V(C_{i+3}))$ in $G(i, i + 3, \Xi, \Psi)$. Thus, $\phi(V(C_{i+1})) \subseteq G(i, i + 3, \Xi, \Psi) \setminus \phi(V(C_{i+3})) = G(i, i + 3, \Xi, \Psi) \setminus (W_i \cup \{w^1_{i+1,\sigma_i}\}) \subseteq W_{i+1} \cup M \cup W_{i+3}$. Since $G[W_{i+1} \cup W_{i+3}]$ is composed of disjoint paths, $\phi(V(C_{i+1}))$ contains at least one vertex from M. Thus, Lemma 6.16 (in Section 6.1.4.2) with $t = k - 1 = i + 1$ gives that $|W_{i+3}| = 2$ and $\phi(V(C_{i+1}))$ contains all of W_{i+3}. Thus, since $|W_{i+3}| = l_{i+3} - 1$ we get that $l_{i+3} = 3$. Since $l_i = l_{i+3}$, we conclude that $l_i = 3$, so $\phi(V(C_i))$ is a 3-cycle in $G(i, i+3, \Xi, \Psi)$. In $P(i, i+3)$, the cycle C_i shares exactly one vertex with C_{i+1} and shares no vertices with C_{i+2} and shares no vertices with C_{i+3}. So, in $G(i, i + 3, \Xi, \Psi)$, $\phi(V(C_i))$ shares exactly one vertex with $\phi(V(C_{i+1}))$ and shares no vertices with either $\phi(V(C_{i+2}))$ or $\phi(V(C_{i+3}))$. Thus, $\phi(V(C_i)) \subseteq G(i, i+3, \Xi, \Psi) \setminus (\phi(V(C_{i+2})) \cup \phi(V(C_{i+3}))) \subseteq G(i, i+3, \Xi, \Psi) \setminus (W_{i+1} \cup W_i) \subseteq M \cup W_{i+3}$. Since $G[W_{i+3}]$ is a path, it follows that $\phi(V(C_i))$ contains at least one vertex from M. However, this contradicts Lemma 6.17 (in Section 6.1.4.2) with $t = i + 3$ and $C = \phi(V(C_{i+1}))$. Since we reached a contradiction under the assumption that $\phi(V(C_{i+2}))$ contains at least two vertices from M, it follows that $\phi(V(C_{i+2}))$ contains exactly one vertex from M, as claimed. Thus, $l_{i+1} = l_{i+2}$, as desired.

Case 3.4.2: Assume that $\phi(V(C_{i+2}))$ contains at least one vertex from W_{i+3}. We will show that this situation is impossible by obtaining a contradiction. Lemma 6.28 (in Section 6.1.4.4) gives that there is $s \in \{i+1, i+3\}$ so that $\phi(V(C_{i+2}))$ contains all of W_s and l_s is odd. We will consider the two cases for s separately.

Case 3.4.2.1: Suppose that $s = i+1$. Then, since $\phi(V(C_{i+2}))$ contains all of W_{i+1} and at least one vertex from W_{i+3}, it follows from the construction of $G(i, i+3, \Xi, \Psi)$ that $\phi(V(C_{i+2}))$ also contains at least one vertex from M. So, it follows that $l_{i+2} = |\phi(V(C_{i+2}))| \geq |W_{i+1}| + 2 = l_{i+1} + 1 > l_{i+1}$. Since $k = i+2$ is defined such that l_k is the smallest odd cycle length that is at least 5, it follows that $l_{i+1} = 3$. So, in $P(i, i+3)$, C_{i+1} is a 3-cycle that shares exactly one vertex with C_{i+2} and shares no vertex with C_{i+3}. Thus, in $G(i, i+3, \Xi, \Psi)$, $\phi(V(C_{i+1}))$ is a 3-cycle that shares exactly one vertex with $\phi(V(C_{i+2}))$ and shares no vertex with $\phi(V(C_{i+3}))$. So, $\phi(V(C_{i+1})) \subseteq G(i, i+3, \Xi, \Psi) \setminus \phi(V(C_{i+3})) = G(i, i+3, \Xi, \Psi) \setminus (W_i \cup \{w^1_{i+1,\sigma_i}\}) \subseteq W_{i+1} \cup M \cup W_{i+3}$. Since $G[W_{i+1} \cup W_{i+3}]$ is composed of disjoint paths, it follows that $\phi(V(C_{i+1}))$ contains at least one vertex from M. Thus, by Lemma 6.16 (in Section 6.1.4.2) with $t = k-1 = i+1$, $|W_{i+3}| = 2$ and $\phi(V(C_{i+1}))$ contains all of W_{i+3}. Thus, since $|W_{i+3}| = l_{i+3} - 1$ we get that $l_{i+3} = 3$, so $\phi(V(C_i))$ is a 3-cycle in $G(i, i+3, \Xi, \Psi)$. In $P(i, i+3)$, the cycle C_i shares exactly one vertex with C_{i+1} and shares no vertices with C_{i+2} and shares no vertices with C_{i+3}. So, in $G(i, i+3, \Xi, \Psi)$, $\phi(V(C_i))$ shares exactly one vertex with $\phi(V(C_{i+1}))$ and shares no vertices with either $\phi(V(C_{i+2}))$

or $\phi(V(C_{i+3}))$. Thus, $\phi(V(C_i)) \subseteq G(i, i+3, \Xi, \Psi) \setminus (\phi(V(C_{i+2})) \cup \phi(V(C_{i+3}))) \subseteq G(i, i+3, \Xi, \Psi) \setminus (W_{i+1} \cup W_i) \subseteq M \cup W_{i+3}$. Since $G[W_{i+3}]$ is a path, it follows that $\phi(V(C_i))$ contains at least one vertex from M. However, this contradicts Lemma 6.17 (in Section 6.1.4.2) with $t = i+3$ and $C = \phi(V(C_{i+1}))$, as desired.

Case 3.4.2.2: Suppose that $s = i+3$. This means that $\phi(V(C_{i+2}))$ contains all of W_{i+3}. However, since $\phi(V(C_{i+2}))$ also contains at least one vertex from W_{i+1} (this was argued at the start of Case 3.4), $\phi(V(C_{i+2}))$ also must contain at least one vertex from M (since M is between W_{i+1} and W_{i+3}, by construction). So, $|\phi(V(C_{i+2}))| \geq |W_{i+3}| + 2$. Thus, since $|\phi(V(C_{i+2}))| = l_{i+2}$ and $|W_{i+3}| = l_{i+3} - 1$, we have that $l_{i+2} > l_{i+3}$ and so, since k was defined as the largest index such that l_k is odd, at least 5, and minimized, it follows that l_{i+3}. We argued at the start of Case 3.4 that $l_i = l_j = l_{i+3}$, so we have shown that $l_i = l_{i+3} = 3$. Since C_i does not share any vertices with cycles C_{i+2} and C_{i+3} in the necklace $P(i, i+3)$, it follows that $\phi(V(C_i))$ is a 3-cycle in $G(i, i+3, \Xi, \Psi)$ that does not share any vertices with $\phi(V(C_{i+2}))$ or $\phi(V(C_{i+3}))$. In particular, this means that $\phi(V(C_i))$ contains no vertices from W_{i+3} (since $\phi(V(C_{i+2}))$ contains all of W_{i+3}) or W_i (we argued at the start of Case 3.4 that $\phi(V(C_j))$ contains all of W_i) and so $\phi(V(C_i)) \subset W_{i+1} \cup M$. Lemma 6.14 and Lemma 6.15 with $s = k-1 = i+1$ imply that the size of the intersection of $\phi(V(C_i))$ and W_{i+1} is neither 0 or 1, so it follows that $\phi(V(C_i))$ contains at least two vertices from W_{i+1}. However, since $G[W_{i+1}]$ is a path and $\phi(V(C_i))$ is a cycle, it follows that $\phi(V(C_i))$ is a 3-cycle consisting of exactly two vertices from W_{i+1} and exactly one vertex from M. By construction, the only two vertices in W_{i+1} that are adjacent to a vertex in M are the endpoints of the path $G[W_{i+1}]$, so it follows that $|W_{i+1}| = 2$ and that $\phi(V(C_i))$ contains all of W_{i+1}, but this is impossible since $\phi(V(C_{i+2}))$ contains at least one vertex from W_{i+1} and $\phi(V(C_i)) \cap \phi(V(C_{i+2})) = \emptyset$. Thus, we have reached the desired contradiction.

Thus, we have shown that $l_{i+1} = l_{i+2}$ in all cases, and we previously showed that $l_i = l_{i+3}$ since $j = i + 3$. Thus, $P(i + 1, i + 3)$ is isomorphic to $P(i, i + 2)$. Now, since $\phi(V(C_{i+3})) = W_i \cup \{w^1_{i+1,\sigma_i}\}$ (by Lemma 6.3), we have that $G(i+1, i+3, \Xi, \Psi)$ contains a copy of $P(i, i+2)$. However, since $P(i+1, i+3)$ is isomorphic to $P(i, i+2)$, this means that $G(i+1, j, \Xi, \Psi)$ contains a copy of $P(i+1, i+3)$ as well. Thus, since $i+1 \leq k \leq i+3$, we apply the induction hypothesis with $i' = i+1$ and $j' = i+3$ to get that $G(k, k, \Xi, \Psi)$ contains a copy of $P(k, k)$, as claimed.

Thus, in all cases, we have shown that $G(k, k, \Xi, \Psi)$ contains a copy of $P(k, k)$, as desired.

6.1.3 Base Cases of Lemma 6.2

The proofs of the base cases largely depend on Lemmas 6.6 and 6.7 below, whose proofs are presented in Sections 6.1.4.4 and 6.1.4.3, respectively. At a high level, these lemmas show that, if there is a copy of C_k that is not contained in M, then the copy of C_k intersects exactly one of W_{k-1} or W_{k+1}, the copy of C_k contains all of whichever of these sets it intersects, and that the lengths of the corresponding cycles are the same.

Lemma 6.6. *Suppose that* $G = G(k-1, k+1, \Xi, \Psi)$ *contains a copy of* $P(k-1, k+1)$ *and let* $\phi : V(P(k-1, k+1)) \to V(G(k-1, k+1, \Xi, \Psi))$ *be the subgraph mapping. Then* $\phi(V(C_k))$ *intersects at most one of* W_{k-1} *or* W_{k+1}*.*

Lemma 6.7. *Let* $\{s, t\} = \{k-1, k+1\}$ *and suppose that one of the following two conditions is true about* $G = G(i, j, \Xi, \Psi)$*.*

1. $G[W_t \cup M]$ *contains a copy of* $P[V(C_t) \cup V(C_k)]$ *via the subgraph mapping* ϕ *so that* $\phi(V(C_k))$ *contains at least one vertex from* W_t*.*

2. $G(k-1, k+1, \Xi, \Psi)$ *contains a copy of* $P(k-1, k+1)$ *via the subgraph mapping* ϕ *so that* $\phi(V(C_k))$ *contains at least one vertex from* W_t *and no vertices from* W_s*.*

Then $\phi(V(C_k))$ *contains all of* W_t *and* $l_t = l_k$*.*

6.1.3.1 Base Case: $j - i = 1$

Suppose that $G(i, j, \Xi, \Psi)$ contains a copy of $P(i, j)$ and let $\phi : V(P(i, j)) \to V(G(i, j, \Xi, \Psi))$ be the subgraph mapping. Our goal is to prove that $G(k, k, \Xi, \Psi)$ contains a copy of $P(k, k)$. Since $j - i = 1$, $P(i, j)$ consists of exactly two cycles C_i and C_j and $k \in \{i, j\}$. Let $t \in \{i, j\}$ be so that $t \neq k$. If $\phi(V(C_k)) \subset M$ then we are done, since $\phi(V(C_k))$ is isomorphic to $C_k = P(k, k)$ and since $G[M] = G(k, k, \Xi, \Psi)$, so suppose that $\phi(V(C_k))$ contains at least one vertex of W_t. Then Condition 1 of Lemma 6.7 is satisfied, so it follows by Lemma 6.7 that $l_t = l_k$, so $\phi(V(C_t))$ forms a copy of C_k in $G(i, j, \Xi, \Psi)$. Moreover, since C_t and C_k are consecutive cycles in the necklace $P(i, j)$, it follows that $\phi(V(C_t))$ shares exactly one vertex with $\phi(V(C_k))$. We claim that $\phi(V(C_t)) \subset M$. Suppose, for the sake of a contradiction, that $\phi(V(C_t))$ contains at least one vertex from W_t. Since it shares exactly one vertex with $\phi(V(C_k))$ and $\phi(V(C_k))$ contains all of W_t (also by Lemma 6.7), this means that $\phi(V(C_t))$ contains exactly one vertex from W_t. But, if $\phi(V(C_t))$ contains exactly one vertex from W_t, then Lemma 6.18 (in Section 6.1.4.3) gives that $|\phi(V(C_t))|$ is even. However, $|\phi(V(C_t))| = l_t = l_k$, which is odd. Thus, we have found a contradiction, and hence $\phi(V(C_t)) \subset M$ as claimed, and so $G(k, k, \Xi, \Psi)$ contains a copy of $P(k, k)$ since $\phi(V(C_t))$ is isomorphic to C_t, which has the same size as $C_k = P(k, k)$, and $G[M] = G(k, k, \Xi, \Psi)$.

6.1.3.2 Base Case: $j - i = 2$ and $k = i + 1$

Suppose that $G(i, j, \Xi, \Psi)$ contains a copy of $P(i, j)$ and let $\phi(V(P(i, j)) \to G(i, j, \Xi, \Psi)$ be the subgraph mapping. Our goal is to prove that $G(k, k, \Xi, \Psi)$ contains a copy of $P(k, k)$. Since $j - i = 2$ and $k = i + 1$, $P(i, j)$ consists exactly of the cycles C_{k-1}, C_k, C_{k+1}. If $\phi(V(C_k))$ is contained in M then we are done, since $\phi(V(C_k))$ is isomorphic to $C_k = P(k, k)$ and since $G[M] = G(k, k, \Xi, \Psi)$, so suppose that it is not. This means that $\phi(V(C_k))$ intersects at least one of W_{k-1} or W_{k+1} and Lemma 6.6 gives that it intersects at most one of W_{k-1} or W_{k+1}, so it intersects exactly one of them. Let $\{s, t\} = \{k - 1, k + 1\}$ be so that $\phi(V(C_k))$ intersects W_t and not W_s. Then Lemma 6.7 gives that $\phi(V(C_k))$ contains all of W_t and $l_t = l_k$, so $\phi(V(C_t))$ forms a copy of C_k in $G(i, j, \Xi, \Psi)$. Moreover, since C_t and C_k are consecutive cycles in the necklace $P(i, j)$, it follows that $\phi(V(C_t))$ shares exactly one vertex with $\phi(V(C_k))$. Since k was defined as the maximum index such that C_k is the smallest odd cycle with length at least 5, it follows that $t \neq k+1$ since $l_t = l_k$. Thus, $t = k-1$. Now, we prove that $\phi(V(C_{k-1})) \subset M$. Suppose, for the sake of a contradiction, that $\phi(V(C_{k-1}))$ contains at least one vertex from W_{k-1} or W_{k+1}. Note that, since $\phi(V(C_k))$ contains all of $W_t = W_{k-1}$ and $\phi(V(C_{k-1}))$ shares exactly one vertex with $\phi(V(C_k))$, it follows that $\phi(V(C_{k-1}))$ contains at most one vertex from W_{k-1}. This gives us two cases: it either contains 0 or 1 vertices from W_{k-1}. We consider these cases separately and show that a contradiction is reached in both cases.

Case 1: $\phi(V(C_{k-1}))$ contains exactly one vertex from W_{k-1}. Now, it is impossible that $\phi(V(C_{k-1}))$ contains no vertices from W_{k+1}, since this would contradict Lemma 6.18 (in Section 6.1.4.3), since we showed above that $l_{k-1} = l_t = l_k$, which is odd. Similarly, it is impossible that it contains exactly one vertex from W_{k+1}, or exactly two vertices from W_{k+1} and no edges from $G[W_{k+1}]$, since these would contradict Lemmas 6.21 and 6.22 (in Section 6.1.4.4), respectively, since we showed above that $l_{k-1} = l_t = l_k$, which is odd. Thus, by Lemma 6.10 (in Section 6.1.4.1), $\phi(V(C_{k-1}))$ contains all of W_{k+1}. This means that $l_k \geq l_{k+1}$. Further, since $\phi(V(C_{k-1}))$ is a cycle with length l_k that contains exactly one vertex from W_{k-1} and contains all of W_{k+1}, Lemma 6.24 (in Section 6.1.4.4) with $C = \phi(V(C_{k-1}))$ and $s = k+1$ tells us that l_{k+1} is odd. Thus, since k is defined as the largest index such that C_k is the smallest odd cycle with size at least 5, it follows that $l_{k+1} = 3$. Thus, $\phi(V(C_{k+1}))$ is a 3-cycle in $G(i, j, \Xi, \Psi)$. Moreover, since C_k and C_{k+1} are consecutive cycles in the necklace $P(i, j)$, it follows that $\phi(V(C_{k+1}))$ shares exactly one vertex with $\phi(V(C_k))$ and no vertices with $\phi(V(C_{k-1}))$. Thus, by Lemma 6.16 (in Section 6.1.4.2), $\phi(V(C_{k+1}))$

contains all of W_{k+1}. However, this is impossible since $\phi(V(C_{k-1}))$ contains all of W_{k+1}, and this would imply that C_{k-1} and C_{k+1}, which are non-consecutive cycles in the necklace $P(i,j)$, are mapped to two cycles that have a non-empty intersection. Thus, we have reached the desired contradiction.

Case 2: $\phi(V(C_{k-1}))$ contains no vertices from W_{k-1}. First, if $\phi(V(C_{k-1}))$ does not contain any vertices from W_{k+1}, then we are done, since this implies that $\phi(V(C_{k-1}))$, which is isomorphic to C_{k-1}, which is isomorphic to $C_k = P(k,k)$ (since $l_{k-1} = l_k$), is contained in M, and $G[M] = G(k,k,\Xi,\Psi)$, so suppose that it contains at least one vertex from W_{k+1}. Now, since $G[W_{k+1}]$ is a path, $\phi(V(C_{k-1}))$ contains at least one vertex from M. Then we apply Lemma 6.20 (in Section 6.1.4.3) with $C = \phi(V(C_{k-1}))$ and $t = k+1$ to get that $\phi(V(C_{k-1}))$ contains all of W_{k+1}. Then since $|W_{k+1}| = l_{k+1} - 1$, this means that $l_k = |\phi(V(C_{k-1}))| \geq l_{k+1}$. By the construction of $G(k-1.k+1,\Xi,\Psi)$, if $\phi(V(C_{k-1}))$ contains all of W_{k+1} and exactly one vertex from M, we have that $l_k = |\phi(V(C_{k-1}))| = |W_{k+1}| + 1 = l_{k+1}$. However, this contradicts the choice of k, since k is defined to be the largest index such that C_k is the smallest odd cycle with size at least 5, so this is impossible. Thus, we find that $\phi(V(C_{k-1}))$ contains at least two vertices from M and $l_k > l_{k+1}$. Together with Lemma 6.20 (in Section 6.1.4.3), we conclude that l_{k+1} is odd. Again, since k is defined to be the largest index such that C_k is the smallest odd cycle with size at least 5, this means that $l_{k+1} = 3$ and so $\phi(V(C_{k+1}))$ is a 3-cycle in $G(i,j,\Xi,\Psi)$. Moreover, since C_k and C_{k+1} are consecutive cycles in the necklace $P(i,j)$, it follows that $\phi(V(C_{k+1}))$ shares exactly one vertex with $\phi(V(C_k))$ and no vertices with $\phi(V(C_{k-1}))$. Then

Lemma 6.16 (in Section 6.1.4.2) gives that $\phi(V(C_{k+1}))$ contains all of W_{k+1}, which is impossible since $\phi(V(C_{k-1}))$ contains all of W_{k+1}, and this would imply that C_{k-1} and C_{k+1}, which are non-consecutive cycles in the necklace $P(i,j)$, are mapped to two cycles that have a non-empty intersection. Thus, we have reached the desired contradiction.

We have found a contradiction in each case, so $\phi(V(C_{k-1})) \subset M$, as claimed. Since $l_{k-1} = l_k$, this means that we have a copy of $C_k = P(k,k)$ contained in $G[M] = G(k,k,\Xi,\Psi)$, as needed.

6.1.4 Technical Lemmas

In this section, we prove the technical results that are necessary for the proof of Lemma 6.2. The relevant definitions of $G(i,j,\Xi,\Psi)$, M, W_t, A, B, k, l_k, etc. can all be found at the start of Section 6.1.

6.1.4.1 General Lemmas

In this section, we make some observations about the structure of the lower bound graph $G(i,j,\Xi,\Psi)$.

Lemma 6.8. *Let C be a cycle of $G = G(i,j,\Xi,\Psi)$. Then C contains a vertex from M if and only if $V(C) \subset W_{k-1} \cup M \cup W_{k+1}$.*

Proof. Suppose that $V(C) \subset W_{k-1} \cup M \cup W_{k+1}$. Then, since $G[W_{k-1} \cup W_{k+1}]$ consists of two disjoint paths, it contains no cycles, so C contains at least one vertex from M.

On the other hand, suppose that C contains a vertex from M. Suppose, for the sake of a contradiction, that C contains a vertex from W_{k-2} or W_{k+2} (since, if it contains a vertex from $V(G) \setminus (W_{k-1} \cup M \cup W_{k+1})$, then it must contain a vertex from one of these sets). We proceed in the case where C contains a vertex from W_{k-2}, the proof if it contains a vertex from W_{k+2} is symmetric. Let $z_1 \in V(C) \cap M$ and $z_2 \in V(C) \cap W_{k-2}$. Then there are two internally disjoint paths in C from z_1 to z_2. However, any path from z_1 to z_2 must contain $w^1_{k-1,\sigma_{k-2}}$ since $w^1_{k-1,\sigma_{k-2}}$ is the vertex in G corresponding to the single vertex shared by cycles C_{k-1} and C_{k-2}. However, by the definition of W_{k-2}, we know that $w^1_{k-1,\sigma_{k-2}} \notin W_{k-2}$, so it follows that $w^1_{k-1,\sigma_{k-2}} \neq z_2$. Similarly, by the definition of M, we know that $w^1_{k-1,\sigma_{k-2}} \notin M$, so $w^1_{k-1,\sigma_{k-2}} \neq z_1$. Thus, $w^1_{k-1,\sigma_{k-2}}$ is an internal vertex in both internally disjoint paths in C from z_1 to z_2, so we have found a contradiction. Thus, $V(C) \subset W_{k-1} \cup M \cup W_{k+1}$, as claimed. □

Corollary 6.9. *Let $t \in \{k-1, k+1\}$. Let C be a cycle of $G = G(i, j, \Xi, \Psi)$ that contains at least one vertex from W_t and at least one vertex from M. Let $z \in V(C) \cap W_t$. Then $N_C(z) \subset W_t \cup M$.*

Proof. Since $N_C(z)$ consists of node z's two neighbours in C, we have $N_C(z) \subseteq V(C)$. Then, by Lemma 6.8, $N_C(z) \subseteq W_{k-1} \cup M \cup W_{k+1}$. However, since no vertex from W_{k-1} is adjacent to a vertex in W_{k+1} by the construction of G, $N_C(z) \subseteq W_t \cup M$, as claimed. □

Lemma 6.10. *Let $t \in \{k-1, k+1\}$. Let C be a cycle of $G = G(i, j, \Xi, \Psi)$ that contains at least one vertex from W_t and at least one vertex from M. Then exactly one of the following is true:*

1. *C intersects W_t exactly once.*

2. *C intersects W_t exactly twice and contains none of the edges of $G[W_t]$.*

3. *C contains all of W_t and all the edges of $G[W_t]$.*

Proof. First, by Lemma 6.8, $V(C) \subseteq W_{k-1} \cup M \cup W_{k+1}$. Next, notice that C intersects W_t either exactly once or at least two times. The former is exactly (1), so consider the latter. Since $G[W_t]$ is a path, this means that C contains both of the endpoints of $G[W_t]$, i.e., $w^1_{t,1}$ and w^1_{t,l_t-1}. Also, since C is a cycle, it contains two internally disjoint paths between $w^1_{t,1}$ and w^1_{t,l_t-1}. It is either the case that neither of these paths is in $G[W_t]$, which corresponds to (2), or that one of these is contained in $G[W_t]$, which corresponds to (3). □

Lemma 6.11. *Let $t \in \{k-1, k+1\}$. Let C be a cycle in $G = G(i, j, \Xi, \Psi)$ of length at most l_k that contains at least one vertex from W_t and at least one vertex from M. Then C does not intersect both A and B.*

Proof. Suppose, for the sake of a contradiction, that C contains a vertex from A and a vertex from B. Since C is a cycle, there exist two internally-disjoint paths P_1 and P_2 in C that each contain exactly one vertex from A and exactly one vertex from B. Recall from the construction of G that M consists of the vertices of $2N$ induced paths in G, where A is one set of endpoints of the paths, and B is the other set of endpoints of the paths. We consider the vertices of these paths, i.e., let $U = \{u^p_q : q \in [l_{k,1}+1] \cup \{0\}, p \in [N]\}$ and $V = \{v^p_q : q \in [l_{k,2}+1] \cup \{0\}, p \in [N]\}$ where $l_k = l_{k,1} + l_{k,2} + 4$, each $u^p_0 = a^p_1$, each $u^p_{l_{k,1}+1} = b^p_1$, each $v^p_0 = a^p_2$, and each $v^p_{l_{k,2}+1} = b^p_2$. Recall from the construction of G that, for index $\sigma_{k,1} = 0$, the edges between vertex sets M and W_{k-1} are precisely $u^p_{\sigma_{k,1}} w^1_{k-1,1}$ and $u^p_{\sigma_{k,1}} w^1_{k-1,l_{k-1}-1}$ for each $p \in [N]$. Moreover, recall from the construction of G that there exists an index $\sigma_{k,2} \in \{0, \dots, l_{k,2} + 1\}$ such that the edges between vertex sets M and W_{k+1} are precisely $v^p_{\sigma_{k,2}} w^1_{k+1,1}$ and $v^p_{\sigma_{k,2}} w^1_{k+1,l_{k+1}-1}$ for each $p \in [N]$. Note that neither P_1 nor P_2 can contain vertices from both U and V, since that would mean that they contain more than one vertex from either A or B. Thus, there are two cases: either P_1 and P_2 both contain vertices from U or V or P_1 contains vertices from U and P_2 contains vertices from V. We consider these cases separately.

Case 1: Let $X \in \{U, V\}$ so that P_1 and P_2 both contain vertices from X. If $X = U$, let $Y = V$, $\epsilon = 1$, and $x = u$, otherwise, let $Y = U$, $\epsilon = 2$, and $x = v$. First, we prove that P_1 and P_2 are disjoint. Suppose, for the sake of contradiction, that P_1 and P_2 are not disjoint. Since they are internally disjoint, this means that they share an endpoint. Let $D \in \{A, B\}$ be so that P_1 and P_2 share an endpoint d in D. Then d has exactly one neighbour in X, at most two neighbours in one of W_s or W_t, where $s \in \{k-1, k+1\} \setminus \{t\}$, and all the remaining neighbours of d are in Y. Then, without loss of generality, the neighbour of d in P_1 is in one of W_s or W_t. However, by the construction of G, this means that P_1 contains another vertex from D, which is a contradiction. Thus, P_1 and P_2 are disjoint. Thus, P_1 contains vertices $a^{p_1}_\epsilon, \dots, x^{p_1}_{\sigma_{k,\epsilon}}, x^{p_2}_{\sigma_{k,\epsilon}+1}, \dots, b^{p_2}_\epsilon$ and P_2 contains vertices $a^{p_3}_\epsilon, \dots, x^{p_3}_{\sigma_{k,\epsilon}}, x^{p_4}_{\sigma_{k,\epsilon}+1}, \dots, b^{p_4}_\epsilon$, where $p_1 \neq p_3$ and $p_2 \neq p_4$. Note that it is possible that $p_1 = p_2$ or $p_3 = p_4$ (or both). Then, since $l_{k,\epsilon} + 2 \geq \lfloor \frac{l_k}{2} \rfloor$ (by the choice of $l_{k,1}$ and $l_{k,2}$) and since we have shown that P_1 and P_2 each contain at least $l_{k,\epsilon} + 2$ vertices, P_1 and P_2 together contain at least $l_k - 1$ vertices. However, since C is a cycle, there is no edge $a_\epsilon^{p_1} a_\epsilon^{p_2}$ and no edge $b_\epsilon^{p_2} b_\epsilon^{p_4}$, so there must be at least two vertices in C that are not in P_1 or P_2. Thus, C contains at least $l_k + 1$ vertices, which is a contradiction.

Case 2: P_1 contains vertices from U and P_2 contains vertices from V. Thus, P_1 and P_2 are disjoint, P_1 is a path $a_1^{p_1}, u_1^{p_1}, \ldots, u_{l_{k,1}}^{p_1}, b_1^{p_1}$ (since $\sigma_{k,1} = 0$), and P_2 contains vertices $a_2^{p_2}, v_1^{p_2}, \ldots, v_{\sigma_{k,2}}^{p_2}, v_{\sigma_{k,2}+1}^{p_3}, \ldots, v_{l_{k,2}}^{p_3}, b_2^{p_3}$, so the total number of vertices on these two paths is $(l_{k,1} + 2) + (l_{k,2} + 2) = l_k$. However, C must contain at least one more vertex: namely, a vertex from W_t. Thus, C contains at least $l_k + 1$ vertices, which is a contradiction.

Thus, since we have found a contradiction in each case, we find that C does not contain a vertex from both A and B, as claimed. □

Corollary 6.12. *Let $t \in \{k-1, k+1\}$. Let C be a cycle in $G = G(i, j, \Xi, \Psi)$ of length at most l_k that contains at least one vertex from W_t and at least one vertex from M. Let $D \in \{A, B\}$ and let $z_1, z_2 \in V(C) \cap D$. Then any path in C from z_1 to z_2 that does not intersect W_{k-1} or W_{k+1} is contained in D.*

Proof. We will prove the statement when $D = A$, as the proof when $D = B$ is symmetric. Lemma 6.8 gives that any path in C from z_1 to z_2 that does not intersect W_{k-1} or W_{k+1} is contained in M. Assume, for the sake of contradiction, that such a path is not contained in A. Then, by the definition of M in the construction of G, the path contains one of the paths $a_q^p, \ldots, b_q^p$ for some $p \in [N]$ and $q \in \{1, 2\}$, so the path contains a vertex from B. However, this contradicts Lemma 6.11, so the path is contained in A, as claimed. □

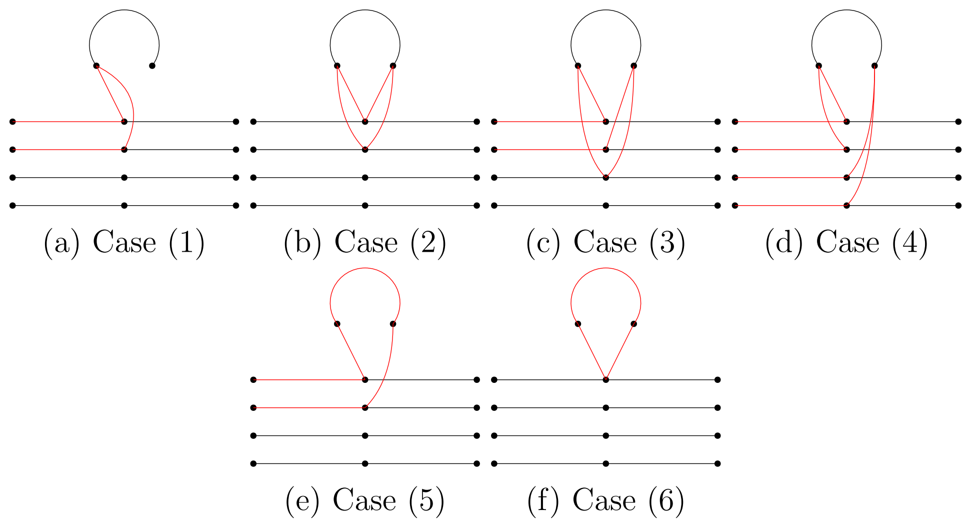

Figure 6.5: The cases of Lemma 6.13. The red edges and paths are those that the lemma tells us are part of C.

Finally, we prove the following lemma, which allows us to be sure that our results in Sections 6.1.4.3 and 6.1.4.4 cover all cases.

Lemma 6.13. *Let $t \in \{k-1, k+1\}$. Let C be a cycle of $G = G(i, j, \Xi, \Psi)$ of length at most l_k that contains at least one vertex from W_t and at least one vertex from M. If $t = k-1$, let $y = u_{\sigma_{k,1}}$ and, if $t = k+1$, let $y = v_{\sigma_{k,2}}$. Then exactly one of the following is true, and we suggest that the reader refer to Figure 6.5 for an illustration of the possibilities.*

(1) There exist distinct vertices $z, x_1, x_2 \in V(C)$ and distinct $p_1, p_2 \in [N]$ such that:

- $V(C) \cap W_t = \{z\}$,
- $N_C(z) = \{y^{p_1}, y^{p_2}\}$,
- *there are two disjoint paths $P_1 = (y^{p_1}, \ldots, x_1)$ and $P_2 = (y^{p_2}, \ldots, x_2)$ in C, and,*

- *there exists $D \in \{A, B\}$ such that both P_1 and P_2 have at least one endpoint in D.*

(2) There exist distinct vertices $z_1, z_2 \in V(C)$ and distinct $p_1, p_2 \in [N]$ such that $V(C) \cap W_t = \{z_1, z_2\}$ and $N_C(z_1) \cup N_C(z_2) = \{y^{p_1}, y^{p_2}\}$. Furthermore, $V(C) = \{z_1, z_2, y^{p_1}, y^{p_2}\}$.

(3) There exist distinct vertices $z_1, z_2 \in V(C)$, distinct vertices $x_1, x_2 \in V(C) \cap M$, and distinct $p_1, p_2, p_3 \in [N]$ such that:

- $V(C) \cap W_t = \{z_1, z_2\}$,
- $y^{p_3} \in V(C) \cap M$,
- $(y^{p_1}, z_1, y^{p_3}, z_2, y^{p_2})$ *is a path in C,*
- *there exist two disjoint paths $P_1 = (y^{p_1}, \dots, x_1)$ and $P_2 = (y^{p_2}, \dots, x_2)$ in $C \cap G[M]$, and,*
- *there exists $D \in \{A, B\}$ such that both P_1 and P_2 have at least one endpoint in D.*

(4) There exist distinct vertices $z_1, z_2 \in V(C)$, distinct vertices $x_1, x_2, x_3, x_4 \in V(C) \cap M$, and distinct $p_1, p_2, p_3, p_4 \in [N]$ such that:

- $V(C) \cap W_t = \{z_1, z_2\}$,
- (y^{p_1}, z_1, y^{p_2}) *and* (y^{p_3}, z_2, y^{p_4}) *are disjoint paths in C,*
- *there exist four disjoint paths $P_1 = (y^{p_1}, \dots, x_1)$, $P_2 = (y^{p_2}, \dots, x_2)$, $P_3 = (y^{p_3}, \dots, x_3)$, and $P_4 = (y^{p_4}, \dots, x_4)$ in $C \cap G[M]$, and,*

- *there exists* $D \in \{A, B\}$ *such that* $P_1, P_2, P_3,$ *and* P_4 *all have at least one endpoint in* D.

(5) There exist distinct vertices $x_1, x_2 \in V(C) \cap M$ *and distinct* $p_1, p_2 \in [N]$ *such that:*

- $W_t \subseteq V(C)$,
- y^{p_1} *is adjacent to* $w^1_{t,1}$ *and* y^{p_2} *is adjacent to* w^1_{t,l_t-1},
- *there exist two disjoint paths* $P_1 = (y^{p_1}, \ldots, x_1)$ *and* $P_2 = (y^{p_2}, \ldots, x_2)$ *in* $C \cap G[M]$, *and,*
- *there exists* $D \in \{A, B\}$ *such that* P_1 *and* P_2 *both have at least one endpoint in* D.

(6) There exists $p \in [N]$ *such that* $W_t \subset V(C)$, $y^p \in V(C) \cap M$, *and* y^p *is adjacent to both* $w^1_{t,1}$ *and* w^1_{t,l_t-1}. *Furthermore,* $V(C) = W_t \cup \{y^p\}$.

Proof. We consider the three cases given by Lemma 6.10 for the intersection between $V(C)$ and W_t.

Case 1: C intersects W_t exactly once. We prove that this condition implies that (1) is satisfied. Let $V(C) \cap W_t = \{z\}$. Now, consider the neighbours of z in C, of which there are exactly two. By Corollary 6.9 and the fact that z is the only vertex of C that is in W_t, it follows that both neighbours of z are in M. By the construction of G, the only edges with one endpoint in M and one endpoint z in W_t are of the form zy^p for some $p \in [N]$. Thus, it follows that $N_C(z) = \{y^{p_1}, y^{p_2}\}$, for some distinct $p_1, p_2 \in [N]$. By the construction of G, y^{p_1} is a vertex in the path $(a^{p_1}, \ldots, b^{p_1})$ of

$G[M]$, so y^{p_1} is the endpoint of a path P_1 in C whose other endpoint is either $a^{p_1} \in A$ or $b^{p_1} \in B$. Similarly, y^{p_2} is a vertex in the path $(a^{p_2}, \ldots, b^{p_2})$ of $G[M]$, so y^{p_2} is the endpoint of a path P_2 in C whose other endpoint is either $a^{p_2} \in A$ or $b^{p_2} \in B$. Moreover, since $p_1 \neq p_2$, the two paths P_1 and P_2 are disjoint in $G[M]$. Finally, by Lemma 6.11, since both paths belong to C, they either both have an endpoint in A or both have an endpoint in B. Thus, the conditions of (1) are satisfied.

Case 2: C intersects W_t exactly twice and contains none of the edges of $G[W_t]$. Let $V(C) \cap W_t = \{z_1, z_2\}$. By Corollary 6.9 and the fact that z_1 and z_2 are the only vertices of C that are in W_t, it follows that $N_C(z_1) \cup N_C(z_2) \subset M$. Now, $\deg_C(z_1) = 2$ and $\deg_C(z_2) = 2$, so $2 \leq |N_C(z_1) \cup N_C(z_2)| \leq 4$. Thus, there are three cases to consider. By the construction of G, the only edges with one endpoint in M and one endpoint z in W_t are of the form zy^p for some $p \in [N]$. Thus, it follows that each vertex in $N_C(z_1) \cup N_C(z_2)$ is a vertex y^p, for some $p \in [N]$.

Case 2.1: $|N_C(z_1) \cup N_C(z_2)| = 2$. We prove that this condition implies that (2) is satisfied. In a cycle C, the fact that $|N_C(z_1) \cup N_C(z_2)| = 2$ implies that $N_C(z_1) = N_C(z_2)$. Thus, there exist distinct $p_1, p_2 \in [N]$ such that $N_C(z_1) = N_C(z_2) = \{y^{p_1}, y^{p_2}\}$. This means that the edges $z_1 y^{p_1}$, $y^{p_1} z_2$, $z_2 y^{p_2}$, and $y^{p_2} z_1$ are in C. However, this means that $V(C) = \{z_1, z_2, y^{p_1}, y^{p_2}\}$, so the conditions of (2) are satisfied.

Case 2.2: $|N_C(z_1) \cup N_C(z_2)| = 3$. We will prove that this condition implies that (3) is satisfied. In the cycle C, each of z_1 and z_2 have exactly two neighbours, so the fact that $|N_C(z_1) \cup N_C(z_2)| = 3$ implies that z_1 and z_2 share exactly one neighbour in C. Thus, there exist distinct $p_1, p_2, p_3 \in [N]$ such that $N_C(z_1) \cup N_C(z_2) = \{y^{p_1}, y^{p_2}, y^{p_3}\}$ with $y^{p_1} \in N_C(z_1)$, $y^{p_2} \in N_C(z_2)$, and $y^{p_3} \in N_C(z_1) \cap N_C(z_2)$. By the construction of G, y^{p_1} is a vertex in the path $(a^{p_1}, \ldots, b^{p_1})$ of $G[M]$, so y^{p_1} is the endpoint of a path P_1 in C whose other endpoint is either $a^{p_1} \in A$ or $b^{p_1} \in B$. Similarly, y^{p_2} is a vertex in the path $(a^{p_2}, \ldots, b^{p_2})$ in $G[M]$, so y^{p_2} is the endpoint of a path P_2 in C whose other endpoint is either $a^{p_2} \in A$ or $b^{p_1} \in B$. Moreover, since $p_1 \neq p_2$, the two paths P_1 and P_2 are disjoint in $G[M]$. Finally, by Lemma 6.11, since P_1 and P_2 are in C, they either both have an endpoint in A or both have an endpoint in B. Thus, the conditions of (3) are satisfied.

Case 2.3: $|N_C(z_1) \cup N_C(z_2)| = 4$. We will show that this condition implies that (4) is satisfied. In the cycle C, each of z_1 and z_2 have exactly two neighbours, so the fact that $|N_C(z_1) \cup N_C(z_2)| = 4$ implies that z_1 and z_2 do not share any neighbours in C. Thus, there exist distinct $p_1, p_2, p_3, p_4 \in [N]$ such that $N_C(z_1) \cup N_C(z_2) = \{y^{p_1}, y^{p_2}, y^{p_3}, y^{p_4}\}$ with $y^{p_1}, y^{p_2} \in N_C(z_1)$ and $y^{p_3}, y^{p_4} \in N_C(z_2)$. By the construction of G, for each $i \in \{1, 2, 3, 4\}$, the vertex y^{p_i} is a vertex in the path $(a^{p_i}, \ldots, b^{p_i})$ of $G[M]$, so y^{p_i} is the endpoint of a path P_i in C whose other endpoint is either $a^{p_i} \in A$ or $b^{p_i} \in B$. Moreover, since the values of p_1, p_2, p_3, p_4 are distinct, the four paths P_1, P_2, P_3, P_4 are disjoint in $G[M]$. Finally, by Lemma 6.11, since each of. the paths P_1, P_2, P_3, P_4 are in C, they either all have an endpoint in A or all have an endpoint in B. Thus, the conditions of (4) are satisfied.

Case 3: C contains all of W_t and all the edges of $G[W_t]$. In particular, this means that C contains the vertices $w^1_{t,1}$ and w^1_{t,l_t-1}, and, since the vertices $w^1_{t,1}$ and w^1_{t,l_t-1} are the endpoints of the path $G[W_t]$, each of these vertices has exactly one neighbour in C that belongs to W_t. Therefore, by Corollary 6.9, $w^1_{t,1}$ and w^1_{t,l_t-1} each have exactly one neighbour in C in M so it follows that $|(N_C(w^1_{t,1}) \cup N_C(w^1_{t,l_t-1})) \cap M| \in \{1, 2\}$. By the construction of G, for each edge that is between a vertex in M and a vertex in W_t, the vertex in W_t is one of the two endpoints of the path $G[W_t]$, so it follows that no vertex other than $w^1_{t,1}$ or w^1_{t,l_t-1}. has a neighbour in M. Moreover, for each edge that is between a vertex in M and a vertex in W_t, the vertex in M is y^p for some $p \in [N]$. Now, if $\left|(N_C(w^1_{t,1}) \cup N_C(w^1_{t,l_t-1})) \cap M\right| = 1$, then the vertices $w^1_{t,1}$ and w^1_{t,l_t-1} have the same neighbour in M, i.e., there exists some $p \in [N]$ such that $y^p \in M$ is adjacent to both $w^1_{t,1}$ and w^1_{t,l_t-1}. It follows that $V(C) = W_t \cup \{y^p\}$, and the conditions of (6) are satisfied. Otherwise, $\left|(N_C(w^1_{t,1}) \cup N_C(w^1_{t,l_t-1})) \cap M\right| = 2$. Then there exist distinct $p_1, p_2 \in [N]$ such that $(N_C(w^1_{t,1}) \cup N_C(w^1_{t,l_t-1})) \cap M = \{y^{p_1}, y^{p_2}\}$ with y^{p_1} adjacent to $w^1_{t,1}$ and with y^{p_2} adjacent to w^1_{t,l_t-1}. By the construction of G, y^{p_1} is a vertex in the path $(a^{p_1}, \dots, b^{p_1})$ in $G[M]$, so y^{p_1} is the endpoint of a path P_1 in C whose other endpoint is either $a^{p_1} \in A$ or $b^{p_1} \in B$. Similarly, y^{p_2} is a vertex in the path $(a^{p_2}, \dots, b^{p_2})$ in $G[M]$, so y^{p_2} is the endpoint of a path P_2 in C whose other endpoint is either $a^{p_2} \in A$ or $b^{p_2} \in B$. Moreover, since $p_1 \neq p_2$, the two paths P_1 and P_2 are disjoint in $G[M]$. Finally, by Lemma 6.11, since P_1 and P_2 are in C, they either both have an endpoint in A or both have an endpoint in B. Thus, the conditions of (5) are satisfied. $\square$

6.1.4.2 Triangles

Throughout our proofs, we find several situations in which we have a subgraph consisting of a cycle (of arbitrary length) and a 3-cycle sharing exactly one vertex, all contained in $W_{k-1} \cup M \cup W_{k+1}$. In this section, we will prove some useful lemmas about this situation. We begin, with Lemmas 6.14 and 6.15, by observing where in $W_{k-1} \cup M \cup W_{k+1}$ a 3-cycle can be.

Lemma 6.14. *There is no 3-cycle with its vertex set contained in M.*

Proof. Suppose, for the sake of contradiction, that there is a 3-cycle with its vertex set contained in M. First, suppose that at least one of the vertices w of the 3-cycle is in $M \setminus (A \cup B)$. In particular, this means that w is a vertex of the form y_q^p where $y \in \{u, v\}$, $p \in [N]$, and either $q \in [l_{k,1}]$ (if $y = u$) or $q \in [l_{k,2}]$ (if $y = v$). However, this leads to a contradiction, since, by the construction of G, such a vertex y_q^p only has two neighbours in G, i.e., the vertices y_{q-1}^p and y_{q+1}^p, but these two vertices are not adjacent in G. Now, the only remaining option is that all three vertices in the cycle belong to $A \cup B$. Recall from the construction of G that the only edges within A are of the form $a_1^{p_1} a_2^{p_2}$ for some $p_1, p_2 \in [N]$, and the only edges within B are of the form $b_1^{p_1} b_2^{p_2}$. First, consider the case in which one vertex belongs to A and another belongs to B. Then there exist $p_1 \in [N]$ and $q_1 \in \{1, 2\}$ so that $a_{q_1}^{p_1} b_{q_1}^{p_1}$ is an edge in the cycle. By the construction of G, the only vertices in $(A \cup B) \setminus \{b_{q_1}^{p_1}\}$ adjacent to $a_{q_1}^{p_1}$ are of the form $a_{q_2}^{p_2}$, where $p_2 \in [N]$ and $q_2 \in \{1, 2\} \setminus \{q_1\}$. However, by the construction of G, no such vertex $a_{q_2}^{p_2}$ is adjacent to $b_{q_1}^{p_1}$. Thus, we have reached a contradiction, so we assume that either all three of the vertices of the cycle are in A or all three of the vertices of the cycle are in B. Thus, the vertices of the 3-cycle are $z_{q_1}^{p_1}, z_{q_2}^{p_2}, z_{q_3}^{p_3}$, where

$z \in \{a,b\}$ and $\{q_1,q_2,q_3\}$ are distinct values from $\{1,2\}$, a contradiction. Thus, we reached a contradiction in all cases, which concludes the proof. □

Lemma 6.15. *Let $s \in \{k-1,k+1\}$. Then there is no 3-cycle with its vertex set contained in $W_s \cup M$ with exactly one vertex in W_s.*

Proof. If $s = k-1$, let $y = u$ and $\epsilon = 1$, and, if $s = k+1$, let $y = v$ and $\epsilon = 2$. Now suppose that there is a 3-cycle with its vertex set contained in $W_s \cup M$ with exactly one vertex in W_s. From the construction of G, all edges between any vertex $w \in W_s$ and any vertex of M is of the form $wy^p_{\sigma_{k,\epsilon}}$ for some $p \in [N]$. This means that the other two vertices in the cycle are $y^{p_1}_{\sigma_{k,\epsilon}}$ and $y^{p_2}_{\sigma_{k,\epsilon}}$, for some distinct $p_1,p_2 \in [N]$. However, by the construction of G, vertices of the form $y^{p_1}_{q_1}$ and $y^{p_2}_{q_2}$ in M are non-adjacent for fixed $y \in \{u,v\}$ whenever $p_1 \neq p_2$, so no such 3-cycle exists. □

Now, we move on to discuss what a 3-cycle sharing exactly one vertex with another cycle can look like in Lemmas 6.16 and 6.17.

Lemma 6.16. *Let $\{s,t\} = \{k-1,k+1\}$. Let C be a cycle in $G = G(i,j,\Xi,\Psi)$ that contains all of W_t and at least one vertex from M. Suppose that there is a 3-cycle C' that intersects M at least once and shares exactly one vertex with C. Then $|W_s| = 2$ and C' contains all of W_s.*

Proof. If $s = k-1$, let $y = u$ and $\epsilon = 1$, and, if $s = k+1$, let $y = v$ and $\epsilon = 2$. First, we will show that the 3-cycle C' does not contain any vertices from W_t. Since C contains all of W_t and C' shares exactly one vertex with C, this means that C' contains at most one vertex from W_t. However, C' cannot contain exactly one vertex from W_t, since this would contradict Lemma 6.15. Now, by Lemma 6.14, C' cannot have its vertex set contained in M, so it must intersect W_s at least once. Further, by Lemma 6.15, C' must intersect W_s at least twice. It is also impossible that all three vertices of C' are contained in W_s, since $G[W_s]$ is a path. Thus, C' contains exactly two vertices from W_s. In particular, this means that the vertices of C' are exactly the two endpoints of the path $G[W_s]$ along with a common neighbour of theirs in M, i.e., the three vertices are $w^1_{s,1}, w^1_{s,l_s-1}$, and $y^p_{\sigma_{k,\epsilon}}$, for some $p \in [N]$. However, since C' is a 3-cycle, this means that $w^1_{s,1}$ and w^1_{s,l_s-1} are adjacent in $G[W_s]$, so $|W_s| = 2$ and C' contains all of W_s, as claimed. □

Lemma 6.17. *Let $t \in \{k-1, k+1\}$ and suppose that $l_t = 3$. Let C be a 3-cycle in $G = G(i, j, \Xi, \Psi)$ that contains at least one vertex from M and all of W_t. Then there is no 3-cycle in G that contains at least one vertex from M and shares exactly one vertex with C.*

Proof. If $t = k-1$, let $y = u$ and $\epsilon = 1$, and, if $t = k+1$, let $y = v$ and $\epsilon = 2$. Suppose, for the sake of contradiction, that there exists a 3-cycle C' in G that contains at least one vertex from M and shares exactly one vertex with C. Defining $s = \{k-1, k+1\} \setminus \{t\}$, Lemma 6.16 gives that C' contains all of W_s. Since $l_t = 3$, it follows that $|W_t| = 2$, so $W_t = \{w^1_{t,1}, w^1_{t,l_t-1}\}$. Since C is a 3-cycle that contains all of W_t and at least one vertex from M, it follows that the vertices of C are exactly $w^1_{t,1}$, w^1_{t,l_t-1}, and $y^p_{\sigma_{k,\epsilon}}$, for some $p \in [N]$. However, this means that C' cannot contain any vertices from W_s, since all the vertices in W_s are at distance at least 2 from all the vertices of C since, if $z = \{u, v\} \setminus \{y\}$, the only vertices in M that are adjacent to a vertex in W_t are of the form z^p_q, i.e., no vertex in W_t is adjacent to $y^p_{\sigma_{k,\epsilon}}$. However, this contradicts the fact that C' contains all of W_t. Thus, there is no such 3-cycle C', which concludes the proof. □

6.1.4.3 Cycles that Intersect One Side

The main goal for this section is to prove Lemma 6.7, though some of the intermediate results will be useful elsewhere. More broadly, the lemmas in this section are concerned with cycles contained in $W_{k-1} \cup M \cup W_{k+1}$ that intersect exactly one of W_{k-1} or W_{k+1}.

Lemma 6.18. *Let $\{s,t\} = \{k-1, k+1\}$. Let C be a cycle in $G = G(i,j,\Xi,\Psi)$ of length at most l_k that contains exactly one vertex from W_t, at least one vertex from M, and no vertices from W_s. Then $|V(C)|$ is even.*

Proof. If $t = k-1$, let $y = u$ and $\epsilon = 1$, and, if $t = k+1$, let $y = v$ and $\epsilon = 2$. Let $z. = V(C) \cap W_t$. By statement (1) of Lemma 6.13, there exist distinct $p_1, p_2 \in [N]$ and $x_1, x_2 \in V(C)$ such that $N_C(z) = \{y^{p_1}_{\sigma_{k,\epsilon}}, y^{p_2}_{\sigma_{k,\epsilon}}$, there exist two disjoint paths $P_1 = (y^{p_1}_{\sigma_{k,\epsilon}}, \dots, x_1)$ and $P_2 = (y^{p_2}_{\sigma_{k,\epsilon}}, \dots, x_2)$ in C, and there exists $D \in \{A, B\}$ such that both P_1 and P_2 have at least one endpoint in D. Suppose, without loss of generality, that they both have an endpoint in A, i.e., that $x_1 = a^{p_1}_\epsilon$ and $x_2 = a^{p_2}_\epsilon$. Then C contains the paths $a^{p_1}_\epsilon, \dots, y^{p_1}_{\sigma_{k,\epsilon}}$ and $a^{p_2}_\epsilon, \dots, y^{p_2}_{\sigma_{k,\epsilon}}$, each having length $\sigma_{k,\epsilon}$. Then the remainder of C is a path between $a^{p_1}_\epsilon$ and $a^{p_2}_\epsilon$, which cannot intersect W_{k-1} or W_{k+1} (by assumption), so Corollary 6.12 gives that it is contained in A. Let r be the number of vertices in this path. Then, since A is bipartite with bipartition $\{a^p_1 : p \in [N]\}, \{a^p_2 : p \in [N]\}$, it follows that r is odd. Now, we have that C contains $1 + 2\sigma_{k,\epsilon} + r$ vertices, which is even. □

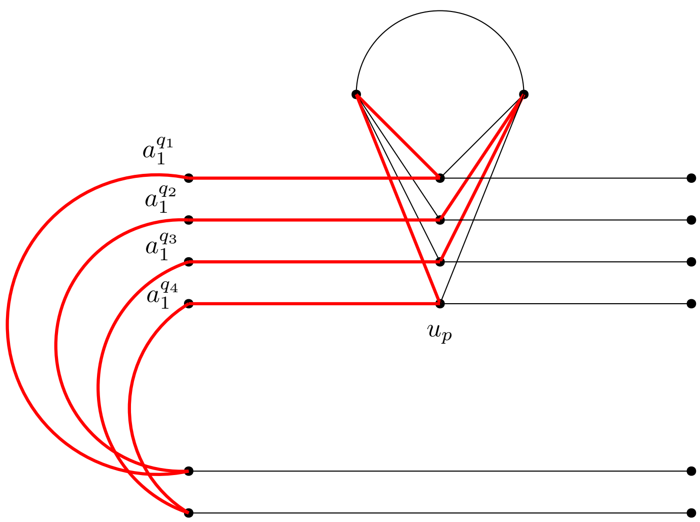

Figure 6.6: An illustration of the proof of case 1 of Lemma 6.19. The red, bold line represents the cycle C. Note that the paths drawn have the smallest possible length, but extending them increases the number of vertices by 2 and hence does not change its parity.

Lemma 6.19. *Let $\{s,t\} = \{k-1, k+1\}$. Let C be a cycle in $G = G(i, j, \Xi, \Psi)$ of length at most l_k that contains exactly two vertices from W_t, at least one vertex from M, no edges from $G[W_t]$, and no vertices from W_s. Then $|V(C)|$ is even.*

Proof. If $t = k-1$, let $y = u$ and $\epsilon = 1$, and, if $t = k+1$, let $y = v$ and $\epsilon = 2$. Let $V(C) \cap W_t = \{z_1, z_2\}$. If C has exactly four vertices then we are done, so the two cases to consider correspond to statements (3) and (4) of Lemma 6.13.

Case 1: Suppose that statement (4) of Lemma 6.13 holds, i.e., that there are distinct vertices $x_1, x_2, x_3, x_4 \in V(C) \cap M$, and distinct $p_1, p_2, p_3, p_4 \in [N]$ such that: $(y^{p_1}_{\sigma_{k,\epsilon}}, z_1, y^{p_2}_{\sigma_{k,\epsilon}})$ and $(y^{p_3}_{\sigma_{k,\epsilon}}, z_2, y^{p_4}_{\sigma_{k,\epsilon}})$ are disjoint paths in C, there exist four disjoint paths $P_1 = (y^{p_1}_{\sigma_{k,\epsilon}}, \ldots, x_1)$, $P_2 = (y^{p_2}_{\sigma_{k,\epsilon}}, \ldots, x_2)$, $P_3 = (y^{p_3}_{\sigma_{k,\epsilon}}, \ldots, x_3)$, and $P_4 = (y^{p_4}_{\sigma_{k,\epsilon}}, \ldots, x_4)$ in $C \cap G[M]$, and, there exists $D \in \{A, B\}$ such that P_1, p_2, P_3, and

P_4 all have at least one endpoint in D. Without loss of generality, assume that all four paths have an endpoint in A, i.e., $x_1 = a_\epsilon^{p_1}$, $x_2 = a_\epsilon^{p_2}$, $x_3 = a_\epsilon^{p_3}$, and $x_4 = a_\epsilon^{p_4}$. Then C contains the following four disjoint paths, each having length $\sigma_{k,\epsilon}$: $a_1^{p_1}, \ldots, y_{\sigma_{k,\epsilon}}^{p_1}$; $a_1^{p_2}, \ldots, y_{\sigma_{k,\epsilon}}^{p_2}$; $a_1^{p_3}, \ldots, y_{\sigma_{k,\epsilon}}^{p_3}$; and $a_1^{p_4}, \ldots, y_{\sigma_{k,\epsilon}}^{p_4}$. Then, without loss of generality, the remainder of C consists of a path from $a_\epsilon^{p_1}$ to $a_\epsilon^{p_2}$ and a path from $a_\epsilon^{p_3}$ to $a_\epsilon^{p_4}$, neither of which can intersect W_{k-1} or W_{k+1} (by assumption). Thus, by Corollary 6.12, these paths are contained in A. Let r_1, r_2 be the number of vertices in these paths, respectively. Then, since A is bipartite with bipartition $\{a_1^p : p \in [N]\}, \{a_2^p : p \in [N]\}$, it follows that r_1 and r_2 are odd. Now, we have that C contains $2 + 4\sigma_{k,\epsilon} + r_1 + r_2$ vertices, which is even.

Case 2: Suppose that statement (3) of Lemma 6.13 holds, i.e., there exist distinct vertices $x_1, x_2 \in V(C) \cap M$, and distinct $p_1, p_2, p_3 \in [N]$ such that: $y_{\sigma_{k,\epsilon}}^{p_3} \in V(C) \cap M$, $(y_{\sigma_{k,\epsilon}}^{p_1}, z_1, y_{\sigma_{k,\epsilon}}^{p_3}, z_2, y_{\sigma_{k,\epsilon}}^{p_2})$ is a path in C, there exist two disjoint paths $P_1 = (y_{\sigma_{k,\epsilon}}^{p_1}, \ldots, x_1)$ and $P_2 = (y_{\sigma_{k,\epsilon}}^{p_2}, \ldots, x_2)$ in $C \cap G[M]$, and, there exists $D \in A, B$ such that both P_1 and P_2 have at least one endpoint in D. Without loss of generality, suppose that P_1 and P_2 both have an endpoint in A, i.e., that $x_1 = a_\epsilon^{p_1}$ and $x_2 = a_\epsilon^{p_2}$. Then C contains the paths $a_\epsilon^{p_1}, \ldots, y_{\sigma_{k,\epsilon}}^{p_1}$ and $a_\epsilon^{p_2}, \ldots, y_{\sigma_{k,\epsilon}}^{p_2}$, each having length $\sigma_{k,\epsilon}$. Then the remainder of C consists of a path from $a_\epsilon^{p_1}$ to $a_\epsilon^{p_2}$ which cannot intersect W_{k-1} or W_{k+1} (by assumption). Thus, Corollary 6.12 gives that it is contained in A. Let r be the number of vertices in this path. Since A is bipartite with bipartition $\{a_1^p : p \in [N]\}, \{a_2^p : p \in [N]\}$, it follows that r is odd. Thus, the number of vertices in C is $2 + 1 + 2\sigma_{k,\epsilon} + r$, which is even. □

Lemma 6.20. *Let $\{s,t\} = \{k-1, k+1\}$. Let C be a cycle in $G = G(i,j,\Xi,\Psi)$ of length l_k that contains at least one vertex from W_t, at least one vertex from M, and no vertices from W_s. Then C contains all of W_t and, if C contains at least two vertices from M, then l_t is odd.*

Proof. First, we show that C contains all of W_t. By Lemma 6.10, it suffices to show that the following two cases are impossible: (1) C contains exactly one vertex from W_t and (2) C contains exactly two vertices from W_t and no edges from $G[W_t]$. Lemma 6.18 implies that (1) cannot be true, since we are assuming that C is a cycle with length exactly l_k, and k was chosen such that l_k is odd. By the same argument, Lemma 6.19 implies that (2) cannot be true. This means that C contains all of W_t.

Next, suppose that C contains at least two vertices from M. If $t = k-1$, let $y = u$ and $\epsilon = 1$, and, if $t = k+1$, let $y = v$ and $\epsilon = 2$. By statement (5) of Lemma 6.13 there exist distinct vertices $x_1, x_2 \in V(C) \cap M$, and distinct $p_1, p_2 \in [N]$ such that: $y^{p_1}_{\sigma_{k,\epsilon}}$ is adjacent to $w^1_{t,1}$ and $y^{p_2}_{\sigma_{k,\epsilon}}$ is adjacent to w^1_{t,l_t-1}, there exist two disjoint paths $P_1 = (y^{p_1}_{\sigma_{k,\epsilon}}, \ldots, x_1)$, $P_2 = (y^{p_2}_{\sigma_{k,\epsilon}}, \ldots, x_2)$ in $C \cap G[M]$, and, there exists $D \in A, B$ such that P_1 and P_2 both have at least one endpoint in D. Without loss of generality, assume that P_1 and P_2 both have at least one endpoint in A, i.e., $x_1 = a^{p_1}_{\epsilon}$ and $x_2 = a^{p_2}_{\epsilon}$. Then C contains the paths $a^{p_1}_{\epsilon}, \ldots, y^{p_1}_{\sigma_{k,\epsilon}}$ and $a^{p_2}_{\epsilon}, \ldots, y^{p_2}_{\sigma_{k,\epsilon}}$. The remainder of C is a path from $a^{p_1}_{\epsilon}$ to $a^{p_2}_{\epsilon}$ that cannot intersect W_{k-1} or W_{k+1} (by assumption). Thus, Corollary 6.12 gives that it is contained in A. Let r be the number of vertices in this path. Then, since A is bipartite with bipartition $\{a^p_1 : p \in [N]\}, \{a^p_2 : p \in [N]\}$, it follows that r is odd. Then the number of vertices in C is $(l_t - 1) + 2\sigma_{k,\epsilon} + r$. Now, the number of vertices in C is l_k, which is odd (by the definition of k), so l_t is odd as well. □

Now, with these results in hand, we are ready to prove Lemma 6.7

Lemma 6.7. *Let* $\{s,t\} = \{k-1, k+1\}$ *and suppose that one of the following two conditions is true about* $G = G(i, j, \Xi, \Psi)$.

1. $G[W_t \cup M]$ *contains a copy of* $P[V(C_t) \cup V(C_k)]$ *via the subgraph mapping* ϕ *so that* $\phi(V(C_k))$ *contains at least one vertex from* W_t.

2. $G(k-1, k+1, \Xi, \Psi)$ *contains a copy of* $P(k-1, k+1)$ *via the subgraph mapping* ϕ *so that* $\phi(V(C_k))$ *contains at least one vertex from* W_t *and no vertices from* W_s.

Then $\phi(V(C_k))$ *contains all of* W_t *and* $l_t = l_k$.

Proof. By Lemma 6.8, $\phi(V(C_k))$ contains at least one vertex from M, so, by Lemma 6.20 with $C = \phi(V(C_k))$, we get that $\phi(V(C_k))$ contains all of W_t.

Now, it is either the case that $\phi(V(C_k))$ contains at least two vertices from M or exactly one vertex from M. We show that it contains exactly one vertex from M. Suppose, for the sake of contradiction, that $\phi(V(C_k))$ contains at least two vertices from M. Then Lemma 6.20 gives that l_t is odd. Moreover, since $\phi(V(C_k))$ contains all of W_t, it follows that $l_k > l_t$. Then since k was chosen so that l_k is the smallest odd length that is at least 5, we must have that $l_t = 3$. Since C_t and C_k are consecutive cycles in the necklace graph P, it follows that $\phi(V(C_t))$ is a 3-cycle in G that shares exactly one vertex with $\phi(V(C_k))$. Now, by the assumptions in the lemma statement, $\phi(V(C_t)) \subseteq W_{k-1} \cup M \cup W_{k+1}$ and, since $G[W_{k-1} \cup W_{k+1}]$ is composed of disjoint paths, $\phi(V(C_t))$ contains at least one vertex from M. Then Lemma 6.16

with $C = \phi(V(C_k))$ and $C' = \phi(V(C_t))$ gives that $\phi(V(C_t))$ contains all of W_s and that $|W_s| = 2$. Now we consider two cases depending on which of the two conditions in the statement of the lemma is true.

Case 1: Suppose that condition (1) is true. Then we have reached a contradiction, since condition (1) assumes that $\phi(V(C_t)) \subseteq W_t \cup M$, but we have just shown that $\phi(V(C_t)) \cap W_s \neq \emptyset$.

Case 2: Suppose that condition (2) is true. Then, since $|W_s| = l_s - 1$, we have that $l_s = 3$. Since C_s and C_k are consecutive cycles in the necklace graph P, it follows that $\phi(V(C_s))$ is a 3-cycle in $G(i, j, \Xi, \Psi)$ that shares exactly one vertex with $\phi(V(C_k))$. Condition (2) assumes that $\phi(V(C_s)) \subseteq W_{k-1} \cup M \cup W_{k+1}$ and, since $G[W_{k-1} \cup W_{k+1}]$ is composed of disjoint paths, $\phi(V(C_s))$ contains at least one vertex from M. Then Lemma 6.16 with $C = \phi(V(C_k))$ and $C' = \phi(V(C_s))$ gives that $\phi(V(C_s))$ contains all of W_s. However, we have separately shown that both $\phi(V(C_t))$ and $\phi(V(C_s))$ contain all of W_s, which means that the intersection of $\phi(V(C_t))$ and $\phi(V(C_s))$ is non-empty, which contradicts the fact that C_t and C_s are not consecutive cycles in the necklace graph P.

Therefore we have derived a contradiction in each case, so $\phi(V(C_k))$ contains exactly one vertex from M, as desired. This, together with the fact that $\phi(V(C_k))$ contains all of W_t, implies that $|\phi(V(C_k))| = |W_t| + 1$. Now, since $|W_t| = l_t - 1$, $|\phi(V(C_k))| = l_k$, and $|\phi(V(C_k))| = |W_t| + 1$, it follows that $l_t = l_k$, as claimed. $\square$

6.1.4.4 Cycles that Intersect Both Sides

The main goal of this section is to prove Lemma 6.6, though some of the intermediate results will be useful elsewhere. More broadly, the lemmas in this section are concerned with cycles contained in $W_{k-1} \cup M \cup W_{k+1}$ that intersect both of W_{k-1} and W_{k+1}.

Lemma 6.21. *Let $\{s,t\} = \{k-1, k+1\}$. Let C be a cycle in $G = G(i,j,\Xi,\Psi)$ of length at most l_k that contains exactly one vertex from W_s and exactly one vertex from W_t. Then $|V(C)|$ is even.*

Proof. Let $z_{k-1} = W_{k-1} \cap V(C)$ and let $z_{k+1} = W_{k+1} \cap V(C)$. By statement (1) of Lemma 6.13 with $t = k-1$, there exist distinct $p_1, p_2 \in [N]$ and $h_1, h_2 \in V(C)$ such that $N_C(z_{k-1}) = u^{p_1}_{\sigma_{k,1}}, u^{p_2}_{\sigma_{k,1}}$, there exist two disjoint paths $P_1 = (u^{p_1}_{\sigma_{k,1}}, \ldots, h_1)$ and $P_2 = (u^{p_2}_{\sigma_{k,1}}, \ldots, h_2)$ in C, and there exists $D \in A, B$ such that both P_1 and P_2 have at least one endpoint in D. Similarly, by statement (1) of Lemma 6.13 with $t = k+1$, there exist distinct $p_3, p_4 \in [N]$ and $h_3, h_4 \in V(C)$ such that $N_C(z_{k+1}) = v^{p_3}_{\sigma_{k,2}}, v^{p_4}_{\sigma_{k,2}}$, there exist two disjoint paths $P_3 = (v^{p_3}_{\sigma_{k,2}}, \ldots, h_3)$ and $P_4 = (v^{p_4}_{\sigma_{k,2}}, \ldots, h_4)$ in C, and there exists $D \in A, B$ such that both P_3 and P_4 have at least one endpoint in D. By Lemma 6.11, since P_1, P_2, P_3, P_4 are all in C, then they all intersect A or they all intersect B. Without loss of generality, assume that they intersect A, i.e., that that $h_1 = a^{p_1}_1$, $h_2 = a^{p_2}_1$, $h_3 = a^{p_3}_2$, and $h_4 = a^{p_4}_2$. Under this assumption, the lengths of paths P_1 and P_2 are each $\sigma_{k,1}$, and the lengths of paths P_3 and P_4 are each $\sigma_{k,2}$. For C to be a cycle, we may assume, without loss of generality, that the rest of the cycle C outside of $z_{k-1}, z_{k+1} \cup P_1 \cup P_2 \cup P_3 \cup P_4$ consists of two disjoint paths: one path P_5 from $a^{p_1}_1$ to $a^{p_3}_2$ and one path P_6 from $a^{p_2}_1$ to $a^{p_4}_2$, neither of which intersect W_{k-1} or

W_{k+1}, so Corollary 6.12 gives that these paths are contained in A. Let r_1 and r_2 be the number of vertices in each of these paths, respectively. Then, since A is bipartite with bipartition $\{a_1^p : p \in [N]\}, \{a_2^p : p \in [N]\}$, it follows that r_1 and r_2 are both even. Thus, the total number of vertices in C is $|\{z_{k-1}, z_{k+1}\}| + |E(P_1)| + |E(P_2)| + |E(P_3)| + |E(P_4)| + |V(P_5)| + |V(P_6)| = 2 + 2\sigma_{k,1} + 2\sigma_{k,2} + r_1 + r_2$, which is even. $\square$

Lemma 6.22. *Let $\{s, t\} = \{k-1, k+1\}$. Let C be a cycle in $G = G(i, j, \Xi, \Psi)$ of length at most l_k that contains exactly one vertex from W_s, exactly two vertices from W_t, and no edges from $G[W_t]$. Then $|V(C)|$ is even.*

Proof. If $s = k-1$, let $y = u$, $x = v$, $\epsilon = 1$, and $\delta = 2$; otherwise, if $s = k+1$, let $y = v$, $x = u$, $\epsilon = 2$, and $\delta = 1$.

First, let $z_s = W_s \cap V(C)$. By statement (1) of Lemma 6.13, there exist distinct $p_1, p_2 \in [N]$ and $h_1, h_2 \in V(C)$ such that $N_C(z_s) = y^{p_1}_{\sigma_{k,\epsilon}}, y^{p_2}_{\sigma_{k,\epsilon}}$, there exist two disjoint paths $P_1 = (y^{p_1}_{\sigma_{k,\epsilon}}, \dots, h_1)$ and $P_2 = (y^{p_2}_{\sigma_{k,\epsilon}}, \dots, h_2)$ in C, and there exists $D \in A, B$ such that both P_1 and P_2 have at least one endpoint in D.

Next, note that it is impossible that C consists of exactly four vertices, since this would mean that exactly one vertex of C is not in W_s or W_t, i.e., there is exactly one vertex in M, but, by the construction of G, there is no vertex in M that is adjacent to a vertex in W_s and to a vertex in W_t. Thus, the only two cases to consider are statements (3) and (4) of Lemma 6.13. In what follows, let $\{z_1, z_2\} = W_t \cap V(C)$.

Case 1: Suppose that statement (4) of Lemma 6.13 holds. Then, there are distinct vertices $h_3, h_4, h_5, h_6 \in V(C) \cap M$, and distinct $p_3, p_4, p_5, p_6 \in [N]$ such that: $(x^{p_3}_{\sigma_{k,\delta}}, z_1, x^{p_4}_{\sigma_{k,\delta}})$ and $(x^{p_5}_{\sigma_{k,\delta}}, z_2, x^{p_6}_{\sigma_{k,\delta}})$ are disjoint paths in C, there exist four disjoint paths $P_3 = (x^{p_3}_{\sigma_{k,\delta}}, \ldots, h_3)$, $P_4 = (x^{p_4}_{\sigma_{k,\delta}}, \ldots, h_4)$, $P_5 = (x^{p_5}_{\sigma_{k,\delta}}, \ldots, h_5)$, and $P_6 = (x^{p_6}_{\sigma_{k,\delta}}, \ldots, h_6)$ in$C \cap G[M]$, and, there exists $D \in A, B$ such that P_3, P_4, P_5, and P_6 all have at least one endpoint in D. By Lemma 6.11, since $P_1, P_2, P_3, P_4, P_5, P_6$ are all in C, then they all intersect A or they all intersect B. Without loss of generality, assume that they intersect A, i.e., that that $h_1 = a^{p_1}_1$, $h_2 = a^{p_2}_1$, $h_3 = a^{p_3}_2$, $h_4 = a^{p_4}_2$, $h_5 = a^{p_5}_2$, and $h_6 = a^{p_6}_2$. Under this assumption, the lengths of paths P_1 and P_2 are each$\sigma_{k,\epsilon}$, and the lengths of paths P_3, P_4, P_5, P_6 are each $\sigma_{k,\delta}$. Since C is a cycle we may assume, without loss of generality, that the rest of C consists of a path from $a^{p_1}_\epsilon$ to $a^{p_3}_\delta$, from $a^{p_2}_\epsilon$ to $a^{p_4}_\delta$, and a path from $a^{p_5}_\delta$ to $a^{p_6}_\delta$, none of which can intersect W_{k-1} or W_{k+1}. By Corollary 6.12, these paths are contained in A. Let r_1, r_2, r_3 be the lengths of these paths, respectively. Then, since A is bipartite with bipartition $\{a^p_1 : p \in [N]\}, \{a^p_2 : p \in [N]\}$, it follows that r_1 and r_2 are even and r_3 is odd. Thus, the total number of vertices in C is $|\{z_s\}| + |\{z_1, z_2\}| + |E(P_1)| + |E(P_2)| + |E(P_3)| + |E(P_4)| + |E(P_5)| + |E(P_6)| + r_1 + r_2 + r_3 = 3 + 2\sigma_{k,\epsilon} + 4\sigma_{k,\delta} + r_1 + r_2 + r_3$, which is even.

Case 2: Suppose that statement (3) of Lemma 6.13 holds, i.e., there exist distinct vertices $h_3, h_4 \in V(C) \cap M$, and distinct $p_3, p_4, p_5 \in [N]$ such that: $x^{p_5}_{\sigma_{k,\delta}} \in V(C) \cap M$, $(x^{p_3}_{\sigma_{k,\delta}}, z_1, x^{p_5}_{\sigma_{k,\delta}}, z_2, x^{p_4}_{\sigma_{k,\delta}})$ is a path in C, there exist two disjoint paths $P_3 = (x^{p_3}_{\sigma_{k,\delta}}, \ldots, h_3)$ and $P_4 = (x^{p_4}_{\sigma_{k,\delta}}, \ldots, h_4)$ in $C \cap G[M]$, and, there exists $D \in A, B$ such that both P_3 and P_4 have at least one endpoint in D. By Lemma 6.11, since P_1, P_2, P_3, P_4 are all in C, then they all intersect A or they all intersect B. Without loss of generality, assume that they intersect A, i.e., that $h_1 = a^{p_1}_1$, $h_2 = a^{p_2}_1$, $h_3 = a^{p_3}_2$, and $h_4 = a^{p_4}_2$. Under this assumption, the lengths of paths P_1 and P_2 are each $\sigma_{k,\epsilon}$, and the lengths of paths P_3 and P_4 are each $\sigma_{k,\delta}$. Since C is a cycle we may assume, without loss of generality, that the rest of C consists of a path from $a_\epsilon^{p_1}$ to $a_\delta^{p_3}$ and a path from $a_\epsilon^{p_2}$ to $a_\delta^{p_4}$, neither of which can intersect W_{k-1} or W_{k+1}. By Corollary 6.12, these paths are contained in A. Let r_1, r_2 be the number of vertices in these paths, respectively. Then, since A is bipartite with bipartition $\{a_1^p : p \in [N]\}, \{a_2^p : p \in [N]\}$, it follows that r_1 and r_2 are even. Thus, the total number of vertices in C is $|\{z_s\}| + |\{z_1, z_2, x_{\sigma_{k,\delta}}^{p_5}\}| + |E(P_1)| + |E(P_2)| + |E(P_3)| + |E(P_4)| + r_1 + r_2 = 4 + 2\sigma_{k,\epsilon} + 2\sigma_{k,\delta} + r_1 + r_2$, which is even. $\square$

Lemma 6.23. *Let $\{s,t\} = \{k-1, k+1\}$. Let C be a cycle in $G = G(i,j,\Xi,\Psi)$ of length at most l_k that contains exactly two vertices from W_{k-1}, no edges from $G[W_{k-1}]$, exactly two vertices from W_{k+1}, and no edges from $G[W_{k+1}]$. Then $|V(C)|$ is even.*

Proof. Let $\{z_1, z_2\} = W_{k-1} \cap V(C)$ and let $\{z_3, z_4\} = W_{k+1} \cap V(C)$. It is impossible that C consists of exactly four vertices, since this would mean that all vertices of C are in W_{k-1} or W_{k+1}, but, by the construction of G, there are no edges between any vertex of W_{k-1} and any vertex of W_{k+1}. Thus, either statement (3) or (4) of Lemma 6.13 applies to each of W_{k-1} and W_{k+1} when considering their relationship to C. We consider each case below.

Case 1: Suppose that statement (4) of Lemma 6.13 applies to both W_{k-1} and W_{k+1}. First, consider W_{k-1}. There are distinct vertices $h_1, h_2, h_3, h_4 \in V(C) \in M$, and distinct $p_1, p_2, p_3, p_4 \in [N]$ such that: the paths $(u_{\sigma_{k,1}}^{p_1}, z_1, u_{\sigma_{k,1}}^{p_2})$ and $(u_{\sigma_{k,1}}^{p_3}, z_2, u_{\sigma_{k,1}}^{p_4})$ are disjoint in C, there exist four disjoint paths $P_1 = (u_{\sigma_{k,1}}^{p_1}, \ldots, h_1)$, $P_2 = (u_{\sigma_{k,1}}^{p_2}, \ldots, h_2)$, $P_3 = (u_{\sigma_{k,1}}^{p_3}, \ldots, h_3)$, $P_4 = (u_{\sigma_{k,1}}^{p_4}, \ldots, h_4)$ in $C \cap G[M]$, and, there exists $D \in \{A, B\}$

such that P_1, P_2, P_3, and P_4 all have at least one endpoint in D.

Similarly, for W_{k+1}, there are distinct vertices $h_5, h_6, h_7, h_8 \in V(C) \cap M$, and distinct $p_5, p_6, p_7, p_8 \in [N]$ such that: the paths $(v^{p_5}_{\sigma_{k,2}}, z_3, v^{p_6}_{\sigma_{k,2}})$ and $(v^{p_7}_{\sigma_{k,2}}, z_4, u^{p_8}_{\sigma_{k,2}})$ are disjoint in C, there exist four disjoint paths $P_5 = (v^{p_5}_{\sigma_{k,2}}, \dots, h_5)$, $P_6 = (v^{p_6}_{\sigma_{k,2}}, \dots, h_6)$, $P_7 = (v^{p_7}_{\sigma_{k,2}}, \dots, h_7)$, $\mathrm{P}_8 = (v^{p_8}_{\sigma_{k,2}}, \dots, h_8)$ in $C \cap G[M]$, and, there exists $D \in \{A, B\}$ such that P_5, P_6, P_7, and P_8 all have at least one endpoint in D.

By Lemma 6.11, since $P_1, P_2, P_3, P_4, P_5, P_6, P_7, P_8$ are all in C, then they all intersect A or they all intersect B. Without loss of generality, assume that they intersect A, i.e., that $h_1 = a^{p_1}_1$, $h_2 = a^{p_2}_1$, $h_3 = a^{p_3}_1$, $h_4 = a^{p_4}_1$, $h_5 = a^{p_5}_2$, $h_6 = a^{p_6}_2$, $h_7 = a^{p_7}_2$, and $h_8 = a^{p_8}_2$. Under this assumption, the lengths of paths P_1, P_2, P_3, P_4 are each $\sigma_{k,1}$, and the lengths of paths P_5, P_6, P_7, P_8 are each $\sigma_{k,2}$. Since C is a cycle, then we can assume without loss of generality that the rest of C consists of a path from $a^{p_1}_1$ to $a^{p_5}_2$, from $a^{p_2}_1$ to $a^{p_6}_2$, from $a^{p_3}_1$ to $a^{p_7}_2$, and from $a^{p_4}_1$ to $a^{p_8}_2$, none of which can intersect W_{k-1} or W_{k+1}. By Corollary 6.12, these paths are contained in A. Let r_1, r_2, r_3, r_4 be the number of vertices in these paths, respectively. Then, since A is bipartite with bipartition $\{a^p_1 : p \in [N]\}, \{a^p_2 : p \in [N]\}$, it follows that r_1, r_2, r_3, and r_4 are all even. Thus, the total number of vertices in C is $|\{z_1, z_2\}| + |\{z_3, z_4\}| + |E(P_1)| + |E(P_2)| + |E(P_3)| + |E(P_4)| + |E(P_5)| + |E(P_6)| + |E(P_7)| + |E(P_8)| + r_1 + r_2 + r_3 + r_4 = 4 + 4\sigma_{k,1} + 4\sigma_{k,2} + r_1 + r_2 + r_3 + r_4$, which is even.

Case 2: Suppose that statement (4) of Lemma 6.13 applies to exactly one of W_{k-1} and W_{k+1}, and that statement (3) applies to the other. Without loss of generality, assume that statement (4) of Lemma 6.13 applies to W_{k+1}. First, consider W_{k+1}, and suppose that statement (4) of Lemma 6.13 holds. Then, there exist distinct vertices $h_1, h_2, h_3, h_4 \in V(C) \cap M$, and distinct $p_1, p_2, p_3, p_4 \in [N]$ such that: the paths $(v^{p_1}_{\sigma_{k,2}}, z_3, v^{p_2}_{\sigma_{k,2}})$ and $(v^{p_3}_{\sigma_{k,2}}, z_4, v^{p_4}_{\sigma_{k,2}})$ are disjoint in C, there exist four disjoint paths $P_1 = (v^{p_1}_{\sigma_{k,2}}, \ldots, h_1)$, $P_2 = (v^{p_2}_{\sigma_{k,2}}, \ldots, h_2)$, $P_3 = (v^{p_3}_{\sigma_{k,2}}, \ldots, h_3)$, $P_4 = (v^{p_4}_{\sigma_{k,2}}, \ldots, h_4)$ in $C \cap G[M]$, and, there exists $D \in \{A, B\}$ such that P_1, P_2, P_3, P_4 all have at least one endpoint in D.

Next, consider W_{k-1}, and suppose that statement (3) of Lemma 6.13 holds. Then, there exist distinct vertices $h_5, h_6 \in V(C) \cap M$, and distinct $p_5, p_6, p_7 \in [N]$ such that: $u^{p_7}_{\sigma_{k,1}} \in V(C) \cap M$, $(u^{p_5}_{\sigma_{k,1}}, z_1, u^{p_7}_{\sigma_{k,1}}, z_2, u^{p_6}_{\sigma_{k,1}})$ is a path in C, there exist two disjoint paths $P_5 = (u^{p_5}_{\sigma_{k,1}}, \ldots, h_5)$ and $P_6 = (u^{p_6}_{\sigma_{k,1}}, \ldots, h_6)$ in$C \cap G[M]$, and, there exists $D \in A, B$ such that both P_5 and P_6 have at least one endpoint in D.

By Lemma 6.11, since $P_1, P_2, P_3, P_4, P_5, P_6$ are all in C, then they all intersect A or they all intersect B. Without loss of generality, assume that they intersect A, i.e., that $h_1 = a^{p_1}_2$, $h_2 = a^{p_2}_2$, $h_3 = a^{p_3}_2$, $h_4 = a^{p_4}_2$, $h_5 = a^{p_5}_1$, $h_6 = a^{p_6}_1$. Under this assumption, the lengths of paths P_1, P_2, P_3, P_4 are each $\sigma_{k,2}$, and the lengths of paths P_5 and P_6 are each $\sigma_{k,1}$. Since C is a cycle, then we can assume without loss of generality that the rest of C consists of a path from $a^{p_1}_1$ to $a^{p_5}_2$, from $a^{p_2}_1$ to $a^{p_6}_2$, and from $a^{p_3}_2$ to $a^{p_4}_2$, none of which can intersect W_{k-1} or W_{k+1}. By Corollary 6.12, these paths are contained in A. Let r_1, r_2, r_3 be the number of vertices in these paths, respectively. Then, since A is bipartite with bipartition $\{a^p_1 : p \in [N]\}, \{a^p_2 : p \in [N]\}$, it follows that r_1 and r_2 are even and r_3 is odd. Thus, the total number of vertices in C is $|\{z_1, z_2\}| + |\{z_3, z_4\}| + |\{u^{p_7}_{\sigma_{k,1}}\}| + |E(P_1)| + |E(P_2)| + |E(P_3)| + |E(P_4)| + |E(P_5)| +$

$|E(P_6)| + r_1 + r_2 + r_3 = 5 + 4\sigma_{k,2} + 2\sigma_{k,1} + r_1 + r_2 + r_3$, which is even.

Case 3: Suppose that statement (3) of Lemma 6.13 applies to both W_{k-1} and W_{k+1}. First, consider W_{k-1}. Then, there exist distinct vertices $h_1, h_2 \in V(C) \cap M$, and distinct $p_1, p_2, p_5 \in [N]$ such that: $u^{p_5}_{\sigma_{k,1}} \in V(C) \cap M$, $(u^{p_1}_{\sigma_{k,1}}, z_1, u^{p_5}_{\sigma_{k,1}}, z_2, u^{p_2}_{\sigma_{k,1}})$ is a path in C, there exist two disjoint paths $P_1 = (u^{p_1}_{\sigma_{k,1}}, \ldots, h_1)$ and $P_2 = (u^{p_2}_{\sigma_{k,1}}, \ldots, h_2)$ in $C \cap G[M]$, and, there exists $D \in A, B$ such that both P_1 and P_2 have at least one endpoint in D.

Next, consider W_{k+1}. Then, there exist distinct vertices $h_3, h_4 \in V(C) \cap M$, and distinct $p_3, p_4, p_6 \in [N]$ such that:$v^{p_6}_{\sigma_{k,2}} \in V(C) \cap M$, $(v^{p_3}_{\sigma_{k,2}}, z_3, v^{p_6}_{\sigma_{k,2}}, z_4, v^{p_4}_{\sigma_{k,2}})$ is a path in C, there exist two disjoint paths $P_3 = (v^{p_3}_{\sigma_{k,2}}, \ldots, h_3)$ and $P_4 = (v^{p_4}_{\sigma_{k,2}}, \ldots, h_4)$ in $C \cap G[M]$, and, there exists $D \in A, B$ such that both P_3 and P_4 have at least one endpoint in D.

By Lemma 6.11, since P_1, P_2, P_3, P_4 are all in C, then they all intersect A or they all intersect B. Without loss of generality, assume that they intersect A, i.e., that that $h_1 = a^{p_1}_1$, $h_2 = a^{p_2}_1$, $h_3 = a^{p_3}_2$, and $h_4 = a^{p_4}_2$. Under this assumption, the lengths of paths P_1 and P_2 are each $\sigma_{k,1}$, and the lengths of paths P_3 and P_4 are each $\sigma_{k,2}$. Since C is a cycle, then we can assume without loss of generality that the rest of C consists of a path from $a^{p_1}_1$ to $a^{p_3}_2$ and from $a^{p_2}_1$ to $a^{p_4}_2$, neither of which can intersect W_{k-1} or W_{k+1}. By Corollary 6.12, these paths are contained in A. Let r_1, r_2 be the number of vertices in these paths, respectively. Then, since A is bipartite, r_1 and r_2 are even. Thus, the total number of vertices in C is $|\{z_1, z_2\}| + |\{z_3, z_4\}| + |\{u^{p_5}_{\sigma_{k,1}}, v^{p_6}_{\sigma_{k,2}}\}| + |E(P_1)| + |E(P_2)| + |E(P_3)| + |E(P_4)| + r_1 + r_2 = 6 + 2\sigma_{k,1} + 2\sigma_{k,2} + r_1 + r_2$, which is even. $\square$

Lemma 6.24. *Let $\{s,t\} = \{k-1, k+1\}$. Let C be a cycle in $G = G(i,j,\Xi,\Psi)$ of length l_k that contains all of W_s, all the edges of $G[W_s]$, and exactly one vertex from W_t. Then l_s is odd.*

Proof. If $s = k+1$, let $y = u$, $x = v$, $\epsilon = 1$, and $\delta = 2$; and, if $s = k-1$, let $y = v$, $x = u$, $\epsilon = 2$, and $\delta = 1$.

First, let $z_t = W_t \cap V(C)$. By statement (1) of Lemma 6.13, there exist distinct $p_1, p_2 \in [N]$ and $h_1, h_2 \in V(C)$ such that $N_C(z_t) = y^{p_1}_{\sigma_{k,\epsilon}}, y^{p_2}_{\sigma_{k,\epsilon}}$, there exist two disjoint paths $P_1 = (y^{p_1}_{\sigma_{k,\epsilon}}, \ldots, h_1)$ and $P_2 = (y^{p_2}_{\sigma_{k,\epsilon}}, \ldots, h_2)$ in C, and there exists $D \in A, B$ such that both P_1 and P_2 have at least one endpoint in D.

Note that it is impossible that C contains one vertex $h \in M$ that is adjacent to both $w^1_{s,1}$ and w^1_{s,l_s-1}, since this would mean that $V(C) = W_s \cup \{h\}$, so $V(C)$ is contained in $W_s \cup M$, i.e., does not intersect W_t. Thus, only statement (5) of Lemma 6.13 applies to W_s. Then, there exist distinct vertices $h_3, h_4 \in V(C) \cap M$ and distinct $p_3, p_4 \cap [N]$ such that: $x^{p_3}_{\sigma_{k,\delta}}$ is adjacent to $w^1_{s,1}$ and $x^{p_4}_{\sigma_{k,\delta}}$ is adjacent to w^1_{s,l_s-1}, there exist two disjoint paths $P_3 = (x^{p_3}_{\sigma_{k,\delta}}, \ldots, h_3)$, $\mathrm{P}_4 = (x^{p_4}_{\sigma_{k,\delta}}, \ldots, h_4)$ in $C \cap G[M]$, and, there exists $D \in A, B$ such that P_3 and P_4 both have at least one endpoint in D.

By Lemma 6.11, since P_1, P_2, P_3, P_4 are all in C, then they all intersect A or they all intersect B. Without loss of generality, assume that they intersect A, i.e., that that $h_1 = a^{p_1}_{\epsilon}$, $h_2 = a^{p_2}_{\epsilon}$, $h_3 = a^{p_3}_{\delta}$, and $h_4 = a^{p_4}_{\delta}$. Under this assumption, the lengths of paths P_1 and P_2 are each $\sigma_{k,\epsilon}$, and the lengths of paths P_3 and P_4 are each $\sigma_{k,\delta}$. Since C is a cycle, then we can assume without loss of generality that the rest of C consists of a path from $a^{p_1}_{\epsilon}$ to $a^{p_3}_{\delta}$ and from $a^{p_2}_{\epsilon}$ to $a^{p_4}_{\delta}$, neither of which can intersect

W_{k-1} or W_{k+1}, so Corollary 6.12 gives that these paths are contained in A. Let r_1 and r_2 be the number of vertices in each of these paths, respectively. Then, since A is bipartite with bipartition $\{a_1^p : p \in [N]\}, \{a_2^p : p \in [N]\}$, it follows that r_1 and r_2 are both even. Thus, the total number of vertices in C is $|\{z_t\}| + |W_s| + |E(P_1)| + |E(P_2)| + |E(P_3)| + |E(P_4)| + r_1 + r_2 = 1 + (l_s - 1) + 2\sigma_{k,\epsilon} + 2\sigma_{k,\delta} + r_1 + r_2$. However, since the total number of vertices in C is l_k, which is odd, it follows that l_s is odd as well. □

Lemma 6.25. *Let $\{s,t\} = \{k-1, k+1\}$. Let C be a cycle in $G = G(i,j,\Xi,\Psi)$ of length l_k that contains all of W_s, all the edges of $G[W_s]$, exactly two vertices from W_t, and no edges from $G[W_t]$. Then l_s is odd.*

Proof. Let $\{z_1, z_2\} = W_t \cap V(C)$. If $s = k+1$, let $y = u$, $x = v$, $\epsilon = 1$, and $\delta = 2$; and, if $s = k-1$, let $y = v$, $x = u$, $\epsilon = 2$, and $\delta = 1$.

Note that it is impossible that C contains one vertex $h \in M$ that is adjacent to both $w_{s,1}^1$ and w_{s,l_s-1}^1, since this would mean that $V(C) = W_s \cup \{h\}$, so $V(C)$ is contained in $W_s \cup M$, i.e., does not intersect W_t. Thus, only statement (5) of Lemma 6.13 applies to W_s. Then, there exist distinct vertices $h_1, h_2 \in V(C) \cap M$ and distinct $p_1, p_2 \in [N]$ such that: $x_{\sigma_{k,\delta}}^{p_1}$ is adjacent to $w_{s,1}^1$ and $x_{\sigma_{k,\delta}}^{p_2}$ is adjacent to w_{s,l_s-1}^1, there exist two disjoint paths $P_1 = (x_{\sigma_{k,\delta}}^{p_1}, \ldots, h_1)$, $P_2 = (x_{\sigma_{k,\delta}}^{p_2}, \ldots, h_2)$ in $C \cap G[M]$, and, there exists $D \in \{A, B\}$ such that P_1 and P_2 both have at least one endpoint in D.

Next, note that it is impossible that C consists of exactly four vertices, since this would mean that all vertices of C are in W_{k-1} or W_{k+1}, but, by the construction of G, there are no edges between any vertex of W_{k-1} and any vertex of W_{k+1}. Thus, either statement (3) or statement (4) of Lemma 6.13 applies to W_t, and we consider these two cases separately below.

Case 1: Suppose that statement (4) of Lemma 6.13 holds. Then, there are distinct vertices $h_3, h_4, h_5, h_6 \in V(C) \cap M$, and distinct $p_3, p_4, p_5, p_6 \in [N]$ such that: $(y^{p_3}_{\sigma_{k,\epsilon}}, z_1, y^{p_4}_{\sigma_{k,\epsilon}})$ and $(y^{p_5}_{\sigma_{k,\epsilon}}, z_2, y^{p_6}_{\sigma_{k,\epsilon}})$ are disjoint paths in C, there exist four disjoint paths $P_3 = (y^{p_3}_{\sigma_{k,\epsilon}}, \ldots, h_3)$, $P_4 = (y^{p_4}_{\sigma_{k,\epsilon}}, \ldots, h_4)$, $P_5 = (y^{p_5}_{\sigma_{k,\epsilon}}, \ldots, h_5)$, and $P_6 = (y^{p_6}_{\sigma_{k,\epsilon}}, \ldots, h_6)$ in $C \cap G[M]$, and, there exists $D \in \{A, B\}$ such that P_3, P_4, P_5, and P_6 all have at least one endpoint in D.

By Lemma 6.11, since $P_1, P_2, P_3, P_4, P_5, P_6$ are all in C, then they all intersect A or they all intersect B. Without loss of generality, assume that they intersect A, i.e., that that $h_1 = a^{p_1}_{\delta}$, $h_2 = a^{p_2}_{\delta}$, $h_3 = a^{p_3}_{\epsilon}$, $h_4 = a^{p_4}_{\epsilon}$, $h_5 = a^{p_5}_{\epsilon}$, and h$_6 = a^{p_6}_{\epsilon}$. Under this assumption, the lengths of paths P_1 and P_2 are each $\sigma_{k,\delta}$, and the lengths of paths P_3, P_4, P_5, P_6 are each $\sigma_{k,\epsilon}$. Since C is a cycle, then we can assume without loss of generality that the rest of C consists of a path from $a^{p_1}_{\delta}$ to $a^{p_5}_{\epsilon}$, from $a^{p_2}_{\delta}$ to $a^{p_1}_{\epsilon}$, and from $a^{p_3}_{\epsilon}$ to $a^{p_4}_{\epsilon}$, none of which can intersect W_{k-1} or W_{k+1}. By Corollary 6.12, these paths are contained in A. Let r_1, r_2, r_3 be the number of vertices in these paths, respectively. Then since A is bipartite with bipartition $\{a^p_1 : p \in [N]\}, \{a^p_2 : p \in [N]\}$, it follows that r_1 and r_2 are even and r_3 is odd. Thus, the total number of vertices in C is $|W_s|+|\{z_1, z_2\}|+|E(P_1)|+|E(P_2)|+|E(P_3)|+|E(P_4)|+|E(P_5)|+|E(P_6)|+r_1+r_2+r_3 = (l_s - 1) + 2 + 2\sigma_{k,\delta} + 4\sigma_{k,\epsilon} + r_1 + r_2 + r_3$. However, since the total number of vertices in C is l_k, which is odd, it follows that l_s is odd as well.

Case 2: Suppose that statement (3) of Lemma 6.13 holds. Then there exist distinct vertices $h_3, h_4 \in V(C) \cap M$, and distinct $p_3, p_4, p_5 \in [N]$ such that: $y^{p_5}_{\sigma_{k,\epsilon}} \in V(C) \cap M$, $(y^{p_3}_{\sigma_{k,\epsilon}}, z_1, y^{p_5}_{\sigma_{k,\epsilon}}, z_2, y^{p_4}_{\sigma_{k,\epsilon}})$ is a path in C, there exist two disjoint paths $P_3 = (y^{p_3}_{\sigma_{k,\epsilon}}, \ldots, h_3)$ and $P_4 = (y^{p_4}_{\sigma_{k,\epsilon}}, \ldots, h_4)$ in $C \cap G[M]$, and, there exists $D \in \{A, B\}$ such that both P_3 and P_4 have at least one endpoint in D.

By Lemma 6.11, since P_1, P_2, P_3, P_4 are all in C, then they all intersect A or they all intersect B. Without loss of generality, assume that they intersect A, i.e., that that $h_1 = a^{p_1}_\delta$, $h_2 = a^{p_2}_\delta$, $h_3 = a^{p_3}_\epsilon$, and $h_4 = a^{p_4}_\epsilon$. Under this assumption, the lengths of paths P_1 and P_2 are each $\sigma_{k,\delta}$, and the lengths of paths P_3 and P_4 are each $\sigma_{k,\epsilon}$. Since C is a cycle, then we can assume without loss of generality that the rest of C consists of a path from $a^{p_1}_\delta$ to $a^{p_3}_\epsilon$ and from $a^{p_2}_\delta$ to $a^{p_4}_\epsilon$, neither of which can intersect W_{k-1} or W_{k+1}. By Corollary 6.12, these paths are contained in A. Let r_1, r_2 be the number of vertices in these paths, respectively. Then since A is bipartite with bipartition $\{a^p_1 : p \in [N]\}, \{a^p_2 : p \in [N]\}$, it follows that r_1 and r_2 are even. Thus, the total number of vertices in C is $|W_s| + |\{z_1, z_2, y^{p_5}_{\sigma_{k,\epsilon}}\}| + |E(P_1)| + |E(P_2)| + |E(P_3)| + |E(P_4)| + r_1 + r_2 = (l_s - 1) + 3 + 2\sigma_{k,\delta} + 2\sigma_{k,\epsilon} + r_1 + r_2$. However, since the total number of vertices in C is l_k, which is odd, l_s is odd as well. □

Lemma 6.26. *Let C be a cycle in $G = G(i, j, \Xi, \Psi)$ of length l_k that contains at least one vertex from W_{k-1} and at least one vertex from W_{k+1}. Then there is $s \in \{k-1, k+1\}$ so that C contains all of W_s and all the edges of $G[W_s]$. Let $t \in \{k-1, k+1\}$ so that $t \neq s$. Then if C does not contain all of W_t and all the edges of $G[W_t]$, then l_s is odd.*

Proof. First, for each $s \in k-1, k+1$, recall that Lemma 6.13 enumerates all possible relationships between W_s and C. If we assume that statements (5) and (6) do not hold for both W_{k-1} and W_{k+1}, then Lemmas 6.21, 6.22, and 6.23 imply that the length of C is even, which contradicts the fact that l_k is odd (by the definition of k). So one of statements (5) or (6) applies to W_s for some $s \in k-1, k+1$, that is, there exists an $s \in k-1, k+1$ such that C contains all of W_s and all the edges of $G[W_s]$.

Next, suppose that C does not contain all of W_t and all the edges of $G[W_t]$. Then Lemma 6.10 gives that it either contains exactly one vertex from W_t or exactly two vertices from W_t and no edges from $G[W_t]$. In the first case, Lemma 6.24 gives that l_s is odd. In the second case, Lemma 6.25 gives that l_s is odd. $\square$

Lemma 6.27. *Let C be a cycle in $G = G(i, j, \Xi, \Psi)$ of length l_k that contains all of W_{k-1}, all of W_{k+1}, all the edges of $G[W_{k-1}]$, and all the edges of $G[W_{k+1}]$. Then exactly one of l_{k-1} or l_{k+1} is odd.*

Proof. Note that it is impossible that C contains one vertex $h \in M$ that is adjacent to both $w^1_{k-1,1}$ and $w^1_{k-1,l_{k-1}-1}$, since this would mean that $V(C) = W_{k-1} \cup \{h\}$, so $V(C)$ is contained in $W_{k-1} \cup M$, i.e., does not intersect W_{k+1}. Using a symmetric argument, it is impossible that C contains one vertex $h \in M$ that is adjacent to both $w^1_{k+1,1}$ and $w^1_{k+1,l_{k+1}-1}$. Thus, statement (5) of Lemma 6.13 applies to both W_{k-1} and W_{k+1}. Applying statement (5) of Lemma 6.13 to W_{k-1}, we get that there exist distinct vertices $h_1, h_2 \in V(C) \cap M$, and distinct $p_1, p_2 \in [N]$ such that $u^{p_1}_{\sigma_{k,1}}$ is adjacent to $w^1_{s,1}$ and $u^{p_2}_{\sigma_{k,1}}$ is adjacent to w^1_{s,l_s-1}, there exist two disjoint paths $P_1 = (u^{p_1}_{\sigma_{k,1}}, \ldots, h_1)$, $P_2 = (u^{p_2}_{\sigma_{k,1}}, \ldots, h_2)$ in $C \cap G[M]$, and, there exists $D \in \{A, B\}$ such that P_1 and P_2 both have at least one endpoint in D.

Applying statement (5) of Lemma 6.13 to W_{k+1}, we get that there exist distinct vertices $h_3, h_4 \in V(C) \cap M$, and distinct $p_3, p_4 \in [N]$ such that $v^{p_3}_{\sigma_{k,2}}$ is adjacent to

$w^1_{t,1}$ and $v^{p_4}_{\sigma_{k,2}}$ is adjacent to w^1_{t,l_t-1}, there exist two disjoint paths $P_3 = (v^{p_3}_{\sigma_{k,2}}, \ldots, h_3)$, $P_4 = (v^{p_4}_{\sigma_{k,2}}, \ldots, h_4)$ in $C \cap G[M]$, and, there exists $D \in \{A, B\}$ such that P_3 and P_4 both have at least one endpoint in D.

By Lemma 6.11, since P_1, P_2, P_3, P_4 are all in C, then they all intersect A or they all intersect B. Without loss of generality, assume that they intersect A, i.e., that that $h_1 = a^{p_1}_1$, $h_2 = a^{p_2}_1$, $h_3 = a^{p_3}_2$, and $h_4 = a^{p_4}_2$. Under this assumption, the lengths of paths P_1 and P_2 are each $\sigma_{k,1}$, and the lengths of paths P_3 and P_4 are each $\sigma_{k,2}$. Since C is a cycle, then we can assume without loss of generality that the rest of C consists of a path from $a^{p_1}_1$ to $a^{p_3}_2$ and from $a^{p_2}_1$ to $a^{p_4}_2$ that cannot intersect W_{k-1} or W_{k+1}, so Corollary 6.12 gives that these paths are contained in A. Let r_1 and r_2 be the number of vertices in each of these paths, respectively. Then, since A is bipartite with bipartition $\{a^p_1 : p \in [N]\}, \{a^p_2 : p \in [N]\}$, it follows that r_1 and r_2 are both even. Thus, the total number of vertices in C is $|W_{k-1}| + |W_{k+1}| + |E(P_1)| + |E(P_2)| + |E(P_3)| + |E(P_4)| + r_1 + r_2 = (l_{k-1} - 1) + (l_{k+1} - 1) + 2\sigma_{k,1} + 2\sigma_{k,2} + r_1 + r_2$. Now, the number of vertices in C is l_k, which is odd, and $2\sigma_{k,1} + 2\sigma_{k,2} + r_1 + r_2 - 2$ is even, so we must have that $l_{k-1} + l_{k+1}$ is odd. Thus, exactly one of l_{k-1} or l_{k+1} is odd, as claimed. $\square$

Lemma 6.28. *Let C be a cycle in $G = G(i, j, \Xi, \Psi)$ of length l_k that contains at least one vertex from W_{k-1} and at least one vertex from W_{k+1}. Then there is $s \in \{k-1, k+1\}$ so that C contains all of W_s and l_s is odd.*

Proof. By Lemma 6.26, there is $s \in \{k-1, k+1\}$ so that C contains all of W_s and all the edges of $G[W_s]$. Let $t \in \{k-1, k+1\}$ so that $t \neq s$. Now, if C does not contain all of W_t and all the edges of $G[W_t]$, then Lemma 6.26 gives that l_s is odd, so we are done. Henceforth, assume that C contains all of W_t and all the edges of $G[W_t]$. Then exactly one of l_s or l_t is odd, by Lemma 6.27. If l_s is even, then swap s and t, and we have that C contains all of W_s and l_s is odd, as claimed. □

Thus, with these results in hand, we are ready to prove Lemma 6.6.

Lemma 6.6. *Suppose that* $G = G(k-1, k+1, \Xi, \Psi)$ *contains a copy of* $P(k-1, k+1)$ *and let* $\phi : V(P(k-1, k+1)) \to V(G(k-1, k+1, \Xi, \Psi))$ *be the subgraph mapping. Then* $\phi(V(C_k))$ *intersects at most one of* W_{k-1} *or* W_{k+1}.

Proof. Suppose, for the sake of a contradiction, that $\phi(V(C_k))$ contains at least one vertex from W_{k-1} and at least one vertex from W_{k+1}. Then Lemma 6.28 gives that there is $s \in \{k-1, k+1\}$ so that $\phi(V(C_k))$ contains all of W_s and l_s is odd. Let $t \in \{k-1, k+1\}$ so that $t \neq s$. Then since $\phi(V(C_k))$ contains at least one vertex from W_s and at least one vertex from W_t, and no vertex in M is adjacent to both a vertex from W_s and a vertex from W_t, it follows from the construction of G that $\phi(V(C_k))$ contains at least two vertices from M. Therefore, $l_k = |\phi(V(C_k))| \geq |W_s| + 2 = (l_s - 1) + 2 = l_s + 1 > l_s$. But, the value of k was defined so that $l_k \geq 5$ is odd and minimized, which implies that $l_s = 3$. By Lemma 6.16 applied to $W_t = W_s$ with $C = \phi(V(C_k))$ and $C' = \phi(V(C_s))$, we get that $\phi(V(C_s))$ contains all of W_t and $|W_t| = 2$. Next, applying Lemma 6.16 to $W_t = W_s$ with $C = \phi(V(C_k))$ and $C' = \phi(V(C_t))$, it follows that $\phi(V(C_t))$ contains all of W_t. However, this is impossible since $\phi(V(C_t)) \cap \phi(V(C_s)) = \emptyset$ (since C_s and C_t are non-adjacent cycles in the necklace P) and $\phi(V(C_s))$ contains all of W_t. Thus we have derived a contradiction, so $\phi(V(C_k))$ intersects at most one of W_{k-1} or W_{k+1}, as claimed. □

6.1.4.5 Bounds on Floors/Ceilings

In this section, we prove the following two numerical identities which we use in our inductive cases.

Lemma 6.29. *For all $m \geq 3$, $\lfloor \frac{m}{2} \rfloor \leq m - 2$.*

Proof. We will show by induction on m.

Base case: $m = 3$. Then $\lfloor \frac{m}{2} \rfloor = 1$ and $m - 2 = 1$, so $\lfloor \frac{m}{2} \rfloor \leq m - 2$.

Induction hypothesis: Let $m \geq 3$ and suppose that $\lfloor \frac{m}{2} \rfloor \leq m - 2$.

Inductive step: Consider $m + 1$. We have that

$$\left\lfloor \frac{m+1}{2} \right\rfloor \leq \left\lfloor \frac{m}{2} \right\rfloor + 1 \leq (m-2) + 1 = (m+1) - 2$$

completing the proof. □

Lemma 6.30. *For all $m \geq 5$, $\lfloor \frac{m}{2} \rfloor \leq m - 3$.*

Proof. We will show by induction on m.

Base case: $m = 5$. Then $\lfloor \frac{m}{2} \rfloor = 2$ and $m - 3 = 2$, so $\lfloor \frac{m}{2} \rfloor \leq m - 3$.

Induction hypothesis: Let $m \geq 5$ and suppose that $\lfloor \frac{m}{2} \rfloor \leq m - 3$.

Inductive step: Consider $m + 1$. We have that

$$\left\lfloor \frac{m+1}{2} \right\rfloor \leq \left\lfloor \frac{m}{2} \right\rfloor + 1 \leq (m-3) + 1 = (m+1) - 3$$

completing the proof. □

6.2 Discussion

Here, we discuss the limitations of Theorem 6.1 in Section 6.1. In particular, we discuss the cases when P contains a triangle or P contains only even cycles, as we do not cover these cases.

First, we discuss why our approach cannot generalize to the case where $l_k = 3$. Drucker et al. [11] provide an explanation for why *any* reduction from two-party set disjointness, in which Alice and Bob each simulate some subset of nodes, does not work for 3-cycles. To briefly restate their argument, they notice that, no matter how the lower bound graph is constructed, one of the players (say Alice) simulates two of the nodes that make up the triangle and the other (Bob, in this case) simulates the third. However, since Alice and Bob must both know all the edges with one endpoint controlled by each of them (i.e. all the cut edges), Alice can immediately detect the triangle without needing any communication.

Next, we discuss the difficulty behind extending our approach to necklaces that contain only even cycles. To do this, we must discuss Drucker et al. [11]'s proof in more detail. In particular, their result is more general than we mentioned in Section 2.1, i.e., they actually show that detecting a cycle C_l takes at least $\Omega(\text{ex}(n, C_l)/(n \log n))$ rounds in the CONGEST model. Their construction is very similar to $G[M]$ in our construction, except that the edges on each "side" of the graph that could be added come from an extremal graph for n, C_l. It is known that, if l is odd, then $\text{ex}(n, C_l) \in \Theta(n^2)$. To see this, notice first that every graph has Túran number at most $O(n^2)$ (since the complete graph on n vertices has $\Theta(n^2)$ edges, and this is the maximum number of edges in an n-vertex graph). Now consider the complete bipartite graph $K_{n/2,n/2}$. Since C_l is not bipartite (since l is odd), then $K_{n/2,n/2}$ does not contain a copy of C_l. Finally, note that $K_{n/2,n/2}$ has $\Theta(n^2)$ edges. Thus, C_l has Túran number $\Theta(n^2)$ with the corresponding extremal graph $K_{n/2,n/2}$. For the case where l is even, it is known that $\mathrm{ex}(n, C_4) \in \Theta(n^{3/2})$ (proved by Kővári et al. [21]), which gives the result that detecting C_4 takes at least $\Omega(\sqrt{n}/\log n)$ in the CONGEST model. Kővári et al. give a method for constructing extremal graphs that realize their asymptotic bound, however, Füredi [15] shows that the exact bound can be realized using Erdős-Rényi polarity graphs. We direct the reader to [28] for a description and construction of these graphs. Thus, there is some hope to construct lower bound graphs for necklace graphs containing 4-cycles based on these extremal graphs and Drucker et al.'s construction, however, we comment that these extremal graphs are not as nice to work with as the complete bipartite graphs that we use for odd cycles and we leave this open for future work.

A natural question to ask is if we can adapt the proof of Korhonen and Rybicki [19] that detecting C_l when $l \geq 6$ is even takes at least $\Omega(\sqrt{n}/\log n)$ rounds in the CONGEST model so that it applies to necklace graphs consisting only of cycles with (even) length at least 6, as their construction does not depend on extremal graphs. At a high level, their lower bound graph differs from the one constructed by Drucker et al. by swapping the roles of edges and paths, i.e., rather than constructing paths and adding edges between paths depending on the input bits to Alice and Bob (as illustrated in Figure 4.1 in Chapter 4), the construction starts with edges and adds paths between endpoints of these edges according to the input bits to Alice and Bob. We show below that our strategy from Section 6.1 will not work, by providing a counterexample.

First, however, we need to describe how our construction would incorporate the Korhonen and Rybicki [19] approach. Unless mentioned otherwise here, the definitions and construction are identical to the construction in the proof from Section 6.1. For the first change, we choose $k \in [m]$ so that $l_k \geq 6$ is even and minimized, and k is maximized. Next, since l_k is even, we have that $l_{k,1} = l_{k,2}$. For convenience, we will write $l' = l_{k,1} = l_{k,2}$. Now, we describe the construction of $P(k,k)$. We begin by adding, for each $p \in [N]$, vertices $a_1^p, b_1^p, a_2^p, b_2^p$ and we add the edges $a_1^p b_1^p$ and $a_2^p b_2^p$. Next, for each $p \in [N^2]$ so that $\Xi_p = 1$, Alice adds vertices $u_1^{(p_1,p_2)}, u_2^{(p_1,p_2)}, \ldots, u_{l'}^{(p_1,p_2)}$ and connects the vertices $a_1^{p_1}, u_1^{(p_1,p_2)}, \ldots, u_{l'}^{(p_1,p_2)}, a_2^{p_2}$ in a path, in the order as written. Similarly, for each $p \in [N^2]$ so that $\Psi_p = 1$, Bob adds vertices $v_1^{(p_1,p_2)}, v_2^{(p_1,p_2)}, \ldots, v_{l'}^{(p_1,p_2)}$ and connects the vertices $b_1^{p_1}, v_1^{(p_1,p_2)}, \ldots, v_{l'}^{(p_1,p_2)}, b_2^{p_2}$ in a path. In the inductive case, for (1) when $i = k$ or (2) when $j = k$, we iterate over all $(p_1, p_2) \in [N] \times [N]$ instead of over all $p \in [N]$. Finally, in the inductive case, we may attempt to add an edge incident to a vertex that does not exist. In this case, we simply skip this edge. Intuitively, this happens when we connect W_{k-1} and W_{k+1} to $G[M]$, since we may not have added all possible "copies" of the paths $a_1^{p_1}, u_1^{(p_1,p_2)}, \ldots, u_{l'}^{(p_1,p_2)}, a_2^{p_2}$ or $b_1^{p_1}, v_1^{(p_1,p_2)}, \ldots, v_{l'}^{(p_1,p_2)}, b_2^{p_2}$. For convenience, we write $A = \{a_q^p : q \in [2], p \in [N]\} \cup \{u_q^{(p_1,p_2)} : q \in [l'], (p_1,p_2) \in [N] \times [N]$ so that $X_{f^{-1}(p_1,p_2)} = 1\}$.

To obtain a counterexample, set P to be the necklace graph composed of the cycles C_8 and C_6. Let $N = 5$. Then, set every bit of Ψ to be 0 and set the bits of Ξ associated with the pairs (1,1), (1,2), (1,4), (1,5) (2,3), (3,4), and (3,5) to 1 and all other bits of Ξ to be 0. Then there is no $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$ but the lower bound graph G contains a copy of P. See Figure 6.7 for an illustration of this. Note that Figure 6.7b only shows part of the lower bound graph P. In particular, it only shows W_1 and A, and, some of the edges and vertex labels have been omitted for clarity.

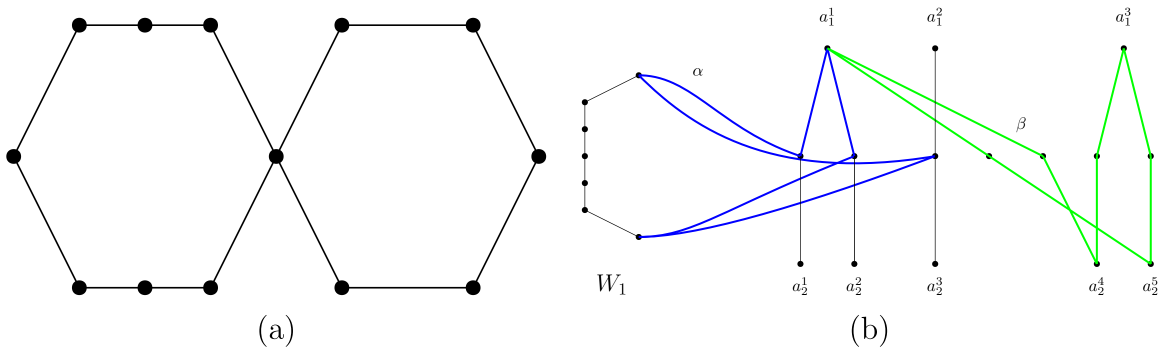

Figure 6.7: (a) Shows the necklace graph P composed of the cycles C_8 and C_6. (b) Shows part of a lower bound graph created for P using our reduction. The cycle C_8 is shown with the bold, green edges (labelled β) and the cycle C_6 is shown using the bold, blue edges (labelled α).

With regards to the above counterexample, a natural question to ask is if we could avoid this situation by maximizing l_k instead of minimizing it, i.e., by choosing $l_k = 8$. We also answer this question in the negative, using the same graph P as above. Also as above, let $N = 5$ and set every bit of Ψ to be 0. Then, set the bits of Ξ associated with the pairs (1,1), (1,2), (2,3), (2,4), and (3,5) to be 1 and all other bits of Ξ to be 0. Then there is no $p \in [N^2]$ so that $\Xi_p = \Psi_p = 1$ but the lower bound graph G contains a copy of P. See Figure 6.8 for an illustration of this. For the sake of clarity, in Figure 6.8, we reversed the order of the cycles, we only show W_1 and A, and we have omitted some edges and vertex labels.

Intuitively, the issue seems to be that the paths added by Alice and Bob based on

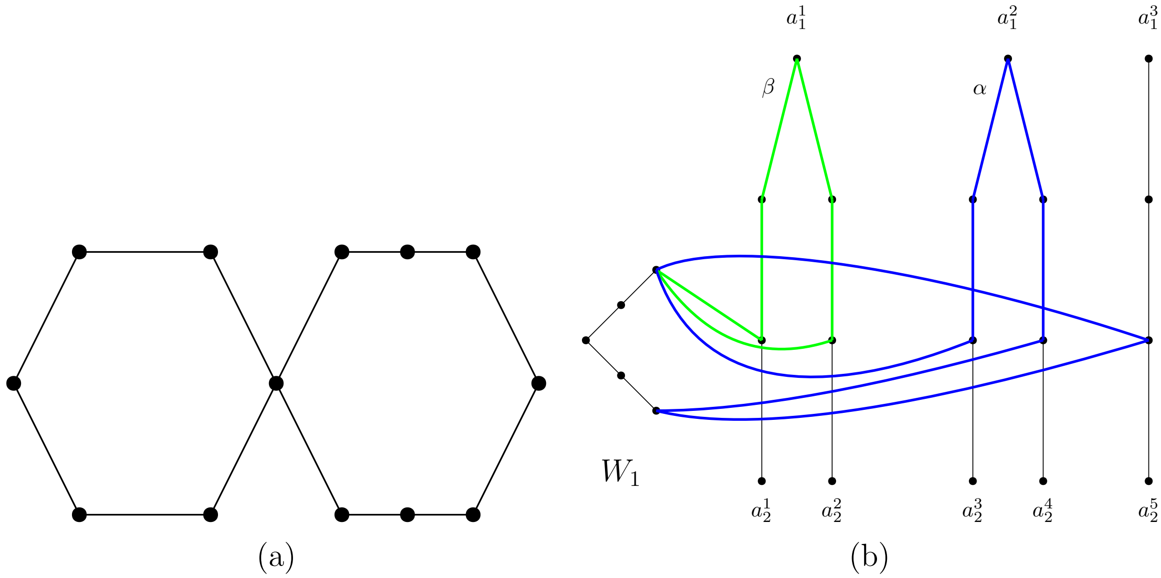

Figure 6.8: (a) shows the necklace graph P composed of the cycles C_6 and C_8. (b) Shows part of a lower bound graph created for P using our reduction. The cycle C_6 is shown with the bold, green edges (labelled β) and the cycle C_8 is shown using the bold, blue edges (labelled α).

their bit strings are significantly less structured than the paths that we used in our proof, which allows many more places in the graph where a cycle could be located. Perhaps there is a way to restrict the construction to remove this possibility.

Chapter 7

Conclusion

We have seen that, if P is a connected graph that contains an induced cycle, then detecting a vertex-rainbow or edge-rainbow copy of P takes at least $\Omega(n/\log n)$ rounds in CONGEST. Further, if P is a necklace graph that contains an odd cycle of length at least 5, we have seen that detecting P takes at least $\Omega(n/\log n)$ rounds in CONGEST. This leaves open the question of what lower bounds could be proved when P is a necklace graph that contains a triangle, or, contains only even cycles. As we discussed in Section 6.2, we believe that it may be possible to adapt our technique when P is a necklace graph containing a 4-cycle, using an extremal graph for n, C_4. It also may be possible to address the problems we have observed when P is a necklace graph containing an even cycle of length at least 6.

Now recall our motivating intuition: If P and Q are graphs so that Q is an induced subgraph of P, then detecting P should be at least as hard as detecting Q. As we have only considered the case in which Q is a cycle, a natural direction for future work is to extend our approach to cases in which Q is not a cycle. Our technique only works if the known lower bound for Q is proved using a reduction from 2-party communication complexity, however, since most lower bounds for subgraph detection are proved using such a reduction, this is not too severe of a limitation. For rainbow subgraph detection, it seems that our techniques can be adapted in a relatively straightforward way, as one can likely still restrict what copies of P in G need to be considered using colours. However, in the uncoloured case, it is not clear if our technique can still be effective if Q is not a cycle, and we leave this question open for future work.

Bibliography

[1] Amir Abboud, Keren Censor-Hillel, Seri Khoury, and Christoph Lenzen. Fooling views: A new lower bound technique for distributed computations under congestion. *Distributed Computing*, 33(6):545–559, 2020.

[2] Zvika Brakerski and Boaz Patt-Shamir. Distributed discovery of large near-cliques. *Distributed Computing*, 24(2):79–89, 2011.

[3] Keren Censor-Hillel, Yi-Jun Chang, François Le Gall, and Dean Leitersdorf. Tight distributed listing of cliques. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2878–2891.

[4] Keren Censor-Hillel, Orr Fischer, Tzlil Gonen, François Le Gall, Dean Leitersdorf, and Rotem Oshman. Fast distributed algorithms for girth, cycles and small subgraphs. *arXiv preprint arXiv:2101.07590*, 2021.

[5] Keren Censor-Hillel, François Le Gall, and Dean Leitersdorf. On distributed listing of cliques. *Proceedings of the 39th Symposium on Principles of Distributed Computing*, pages 474–482, Jul 2020.

[6] Keren Censor-Hillel, Petteri Kaski, Janne H. Korhonen, Christoph Lenzen, Ami

Paz, and Jukka Suomela. Algebraic methods in the congested clique. *Distributed Computing*, 32(6):461–478, Mar 2016.

[7] Yi-Jun Chang, Seth Pettie, and Hengjie Zhang. Distributed triangle detection via expander decomposition. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 821–840. SIAM, 2019.

[8] Yi-Jun Chang and Thatchaphol Saranurak. Improved distributed expander decomposition and nearly optimal triangle enumeration. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, PODC '19, page 66–73, New York, NY, USA, 2019. Association for Computing Machinery.

[9] Artur Czumaj and Christian Konrad. Detecting cliques in congest networks. *Distributed Computing*, 33:533–543, Dec 2020.

[10] Danny Dolev, Christoph Lenzen, and Shir Peled. "tri, tri again": finding triangles and small subgraphs in a distributed setting. In *International Symposium on Distributed Computing*, pages 195–209. Springer, 2012.

[11] Andrew Drucker, Fabian Kuhn, and Rotem Oshman. On the power of the congested clique model. In *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing*, PODC '14, page 367–376, New York, NY, USA, 2014. Association for Computing Machinery.

[12] Talya Eden, Nimrod Fiat, Orr Fischer, Fabian Kuhn, and Rotem Oshman. Sublinear-Time Distributed Algorithms for Detecting Small Cliques and Even Cycles. In Jukka Suomela, editor, *33rd International Symposium on Dis-*

tributed Computing (DISC 2019), volume 146 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 15:1–15:16, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[13] Orr Fischer, Tzlil Gonen, Fabian Kuhn, and Rotem Oshman. Possibilities and impossibilities for distributed subgraph detection. SPAA '18, page 153–162, New York, NY, USA, 2018. Association for Computing Machinery.

[14] Orr Fischer, Tzlil Gonen, and Rotem Oshman. Superlinear lower bounds for distributed subgraph detection. *arXiv preprint arXiv:1711.06920*, 2017.

[15] Zoltán Füredi. Quadrilateral-free graphs with maximum number of edges. In *Proceedings of the Japan Workshop on Graph Th. and Combinatorics*, pages 13–22, 1994.

[16] Mohsen Ghaffari, Fabian Kuhn, and Hsin-Hao Su. Distributed mst and routing in almost mixing time. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, PODC '17, page 131–140, New York, NY, USA, 2017. Association for Computing Machinery.

[17] Mohsen Ghaffari and Jason Li. New distributed algorithms in almost mixing time via transformations from parallel algorithms. *arXiv preprint arXiv:1805.04764*, 2018.

[18] Taisuke Izumi and François Le Gall. Triangle finding and listing in congest networks. *Proceedings of the ACM Symposium on Principles of Distributed Computing*, pages 381–389, Jul 2017.

[19] Janne H Korhonen and Joel Rybicki. Deterministic subgraph detection in broadcast congest. *arXiv preprint arXiv:1705.10195*, 2017.

[20] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.

[21] Tamás Kővári, Vera T. Sós, and Paul Turán. On a problem of k. zarankiewicz. *Colloquium Mathematicum*, 3:50–57, 1954.

[22] François Le Gall and Masayuki Miyamoto. Lower Bounds for Induced Cycle Detection in Distributed Computing. In Hee-Kap Ahn and Kunihiko Sadakane, editors, *32nd International Symposium on Algorithms and Computation (ISAAC 2021)*, volume 212 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 58:1–58:19, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[23] Christoph Lenzen. Optimal deterministic routing and sorting on the congested clique. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing*, PODC '13, page 42–50, New York, NY, USA, 2013. Association for Computing Machinery.

[24] Christoph Lenzen and Roger Wattenhofer. Tight bounds for parallel randomized load balancing. *Distributed Computing*, 29(2):127–142, 2016.

[25] Amir Nikabadi and Janne Korhonen. Beyond distributed subgraph detection: Induced subgraphs, multicolored problems and graph parameters. In *25th International Conference on Principles of Distributed Systems*, volume 217, 2022.

[26] Gopal Pandurangan, Peter Robinson, and Michele Scquizzato. On the distributed complexity of large-scale graph computations. *ACM Trans. Parallel Comput.*, 8(2), jul 2021.

[27] Boaz Patt-Shamir and Marat Teplitsky. The round complexity of distributed sorting: Extended abstract. In *Proceedings of the 30th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, PODC '11, page 249–256, New York, NY, USA, 2011. Association for Computing Machinery.

[28] Xing Peng, Mike Tait, and Craig Timmons. On the chromatic number of the erdos-rényi orthogonal polarity graph. *arXiv preprint arXiv:1408.4065*, 2014.