UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

TOTALLY REFLECTED GROUPS

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

RACHEL WRIGHT
Norman, Oklahoma
2016

TOTALLY REFLECTED GROUPS

A DISSERTATION APPROVED FOR THE
DEPARTMENT OF MATHEMATICS

BY

Dr. Andrew Miller, Chair

Dr. Murad Ozaydin

Dr. Leonard Rubin

Dr. Sepideh Stewart

Dr. Stacy Reeder

Acknowledgments

I will forever be indebted to my advisor, Dr. Miller, who gave generously of his time and knowledge. My work here would not have been possible without his guidance and encouragement. His patience is without bound.

I am thankful for my husband Ty for convincing me that I was cut out for graduate school and for believing in me even in those moments when I did not believe in myself. He helped me see that I could be a mathematician when I grew up.

I cannot thank my sister and best friend Lauren enough for all the love and support she has given me every day of my life. There is a strong possibility that my love of math is due to her teaching me (for fun) how to add and subtract on the blackboards in our mother's classroom.

I would like to thank Drs. Rubin, Stewart, Ozaydin, and Reeder for their participation on my committee.

I am much obliged to my classmate Cherith who, in going before me, made the path clearer.

I am grateful for the many wonderful math teachers I had growing up and attending school in Clinton, Oklahoma, especially Mrs. Donna Lewallen, Mrs. Deb Morris, and Mrs. Jeanette Stone. I would also like to thank my high school teachers Mrs. Janice Berrong and Coach Clay Chittum for teaching me how to study and work hard.

I give thanks daily for my parents for more reasons than I can mention here. They instilled in me a love of learning and a desire to do my very best in all things.

Table of Contents

List of Figures

Abstract

A group G is totally reflected if it has a generating set $\mathcal{S}$ such that each edge in the Cayley graph $\Gamma = \Gamma(G, \mathcal{S})$ is inverted by some color-preserving graph reflection on Γ. For example, we will show that Coxeter groups and right-angled Artin groups are totally reflected and that a finitely generated abelian group is totally reflected if and only if its first invariant factor is even. We show that direct and free products of totally reflected groups are totally reflected. More generally, we develop a group construction called a right-angled product which generalizes free and direct products, and we show that a right-angled product of totally reflected groups is itself totally reflected.

A group G is strongly totally reflected if there exists a color-preserving reflection group G_R acting on $\Gamma(G, \mathcal{S})$ such that each edge in the graph is inverted by some reflection in G_R. We state and prove sufficient conditions for a totally reflected group to be strongly totally reflected and use these results to prove from a graphical perspective that any right-angled Artin group is commensurable with a right-angled Coxeter group. In particular, we show that both the right-angled Artin group $A_\Delta = \langle \mathcal{S} \rangle$ and its associated right-angled Coxeter group A_r are finite-index subgroups of the group of color-preserving graph automorphisms of $\Gamma(A_\Delta, \mathcal{S})$.

Chapter 1

Introduction

Let's begin our journey in the familiar territory of 2-dimensional Euclidean space, $\mathbb{R}^2$. Consider the lines l_1, l_2 and l_3 given by the equations $y = 0$, $y = \sqrt{3} \cdot x$, and $y = -\sqrt{3} \cdot x$, respectively. The smallest angle between any two of these lines is exactly $\pi/3$. Each line serves as the axis for a reflection in $\mathbb{R}^2$. Let r_1, r_2, and r_3 denote the automorphisms of $\mathbb{R}^2$ corresponding to the reflections in the lines l_1, l_2, and l_3, respectively. The reflections are isometries of the real plane, and the composition of two isometries is again an isometry. Therefore, r_1, r_2, and r_3 generate a group, say G, with binary operation given by composition.

It is well known that the composition of two distinct reflections in $\mathbb{R}^2$ whose axes intersect at a single point P is equivalent to a rotation about P. In our specific example, the compositions r_2r_1 (meaning, reflection across l_1 followed by reflection across l_2), r_3r_2, and r_1r_3 are each equivalent to counterclockwise rotation about the origin through an angle with measure $2\pi/3$, while r_1r_2, r_2r_3, and r_3r_1 are each equivalent to clockwise rotation about the origin through an angle with measure $2\pi/3$. Figure 1.1 illustrates the situation where r_1 is then followed by r_2.

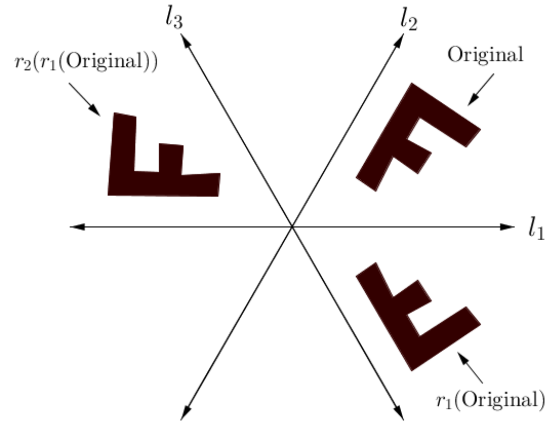

Figure 1.1: Rotational effect of composing reflections in $\mathbb{R}^2$

We can observe that reflection in the line l_3 can be achieved as a composition of reflections in the lines l_1 and l_2. Specifically, $r_3 = r_1 r_2 r_1 = r_2 r_1 r_2$. Therefore, the group G can be generated by r_1 and r_2 alone. One can check that G has exactly six elements, namely 1_G (identity), r_1, r_2, $r_1 r_2 r_1$, $r_2 r_1$, and $r_1 r_2$. Of course, there may be other ways to write these elements in terms of r_1 and r_2.

If we carefully arrange an equilateral triangle, T, in $\mathbb{R}^2$ so that its centroid is at the origin, in the manner depicted in Figure 1.2, we can see that the lines l_1, l_2, and l_3 are precisely the three lines of reflectional symmetry for the triangle. Therefore, the reflections r_1, r_2, and r_3 leave T invariant. In addition to the three reflectional symmetries, there are also three distinct (counterclockwise) rotational symmetries on the triangle. We will denote these as R_0, R_{120}, and R_{240}, where the subscript corresponds to the degree measure of counterclockwise rotation about the origin. The set $\{r_1, r_2, r_3, R_0, R_{120}, R_{240}\}$ is the complete set of symmetries on T, and it forms a group with the binary operation of composition. This group is commonly referred to as the dihedral group with six elements and denoted by the symbol D_3. It is well known that D_3 can be generated by the elements r_1 and r_2 alone. Therefore, $D_3 = \langle r_1, r_2 \rangle = G$. The element R_0 corresponds to the identity element 1_G of the group. The rotations R_{120} and R_{240} can be written in terms of r_1 and r_2. Specifically, $R_{120} = r_2 r_1$ and $R_{240} = r_1 r_2 = (r_2 r_1)^2$.

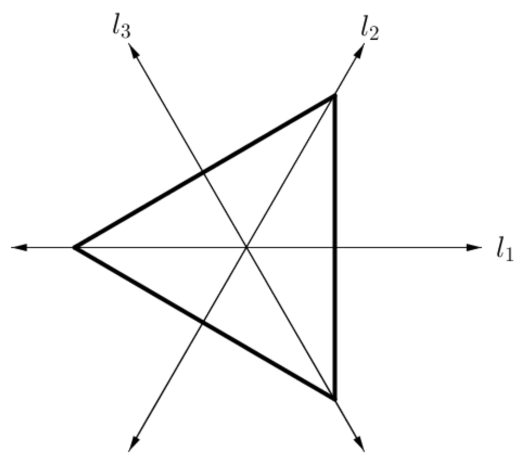

Figure 1.2: Lines of symmetry for the equilateral triangle, T

Observe that the subset $\mathbb{R}^2 - (l_1 \cup l_2 \cup l_3)$ of $\mathbb{R}^2$ consists of six disjoint, open, convex regions, called *chambers*. We will label the chambers with the symbols C_i for each $i \in \{1,2,3,4,5,6\}$, as depicted in Figure 1.3. There are six pairs of chambers which we can think of as being *adjacent*. For example, C_1 and C_2 are adjacent, but C_1 and C_3 are not. In other words, two chambers are adjacent if their closures in $\mathbb{R}^2$ with respect to the standard topology intersect in a ray based at the origin.

We can observe that the group $G = D_3 = \langle r_1, r_2 \rangle$ acts on $\mathbb{R}^2$. Moreover, this action is simply transitive on the set of chambers $\{C_1, C_2, C_3, C_4, C_5, C_6\}$. Choose a point P_1 in the chamber C_1 which is equidistant from lines l_1 and

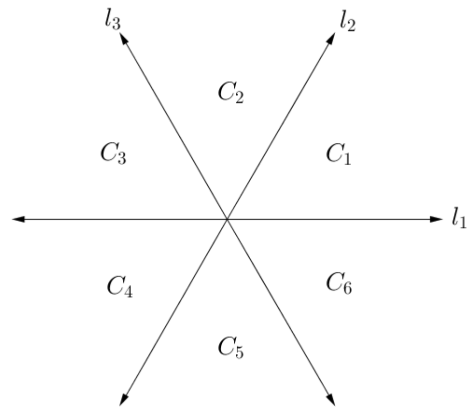

Figure 1.3: Chambers of $\mathbb{R}^2 - (l_1 \cup l_2 \cup l_3)$

l_2. Suppose that we draw a simple closed curve $\mathcal{C}$ in $\mathbb{R}^2$ as described here: Place a single point at each element of the orbit of P_1 under the action of G and connect two points by a line segment if the chambers which contain them are adjacent. The curve $\mathcal{C}$ is depicted in Figure 1.4. We will informally call it the *chamber adjacency graph for the action of* G *on the plane.* Notice that because of the way the curve $\mathcal{C}$ is situated in $\mathbb{R}^2$, the reflections r_1, r_2, and r_3 of $\mathbb{R}^2$ act as reflections on $\mathcal{C}$ as well.

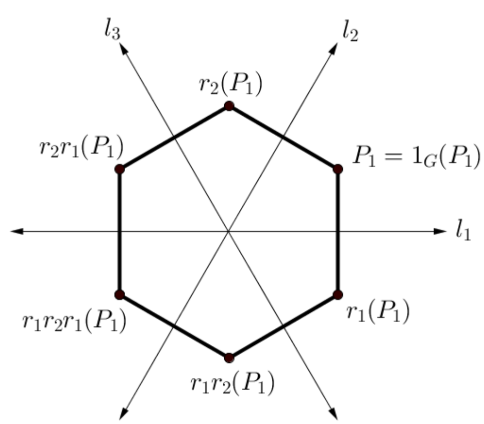

Figure 1.4: Chamber adjacency graph for the action of G on $\mathbb{R}^2$

Some readers may recognize $\mathcal{C}$ as the Cayley graph $\Gamma = \Gamma(D_3, \{r_1, r_2\})$.

In general, the vertices in a Cayley graph $\Gamma(G,\mathcal{S})$ correspond to, and are labeled by, the elements of the group G. The edges in the graph $\Gamma(G,\mathcal{S})$ are *directed* and *colored* line segments or curves. An edge's *color* is an element of $\mathcal{S}$. All of the edges in $\Gamma(G,\mathcal{S})$ which are colored by the same element $s \in \mathcal{S}$ are drawn using the same style (solid or dashed, for instance) or the same color. Vertices labeled by the group elements g and h are connected by an s-colored edge with direction arrow pointing from g to h whenever $g^{-1}h = s$. Said differently, if we begin at a vertex g and travel across an s-colored edge in the direction of the arrow and arrive at a vertex h, then $h = gs$.

The colored and directed Cayley graph $\Gamma = \Gamma(D_3, \{r_1, r_2\})$ is shown in Figure 1.5. The solid edges are those colored by the generator r_1 and the dashed edges are those colored by the generator r_2. As we have done in Figure 1.5, it is conventional to leave off the direction arrow when the color of the edge corresponds to a generator of order 2 in the group. This makes sense because if s has order 2, then $h = gs$ and $g = hs$.

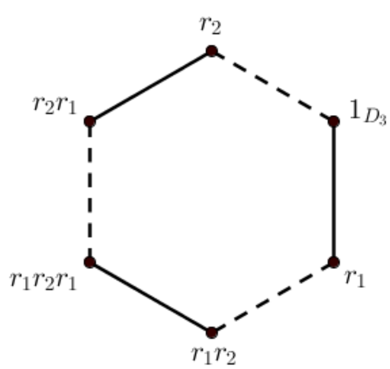

Figure 1.5: Cayley graph $\Gamma(D_3, \{r_1, r_2\})$

We will show that each element g of D_3 corresponds to a *graph automorphism* L_g on Γ, where L_g corresponds to *left translation* (or left multiplication). For example, L_{r_1} sends a vertex labeled by an element $g \in D_3$ to the vertex labeled by $r_1 g \in D_3$. Since r_1 has order 2 in the group, the graph automorphism L_{r_1} has order 2 also. In other words, $(L_{r_1})^2$ fixes the entire graph Γ. One can easily check that L_{r_1} interchanges the vertices 1_{D_3} and r_1, r_2 and $r_1 r_2$, and $r_2 r_1$ and $r_1 r_2 r_1$. This action also preserves vertex adjacencies. That is, if vertices labeled by group elements g and h are connected by an edge in Γ, then the vertices labeled by the elements $r_1 g$ and $r_1 h$ are connected by an edge in Γ. The action of L_{r_1} on Γ *inverts* the edges between 1_{D_3} and r_1 and between $r_2 r_1$ and $r_1 r_2 r_1$. These inverted edges *separate* the graph Γ. Therefore, we will say that L_{r_1} is a *graph reflection* on Γ. The graph automorphisms L_{r_2} and L_{r_3} are also graph reflections on Γ

Suppose that we now consider the Cayley graph for the group D_3 with respect to the generating set $\mathcal{T} = \{a, b\}$, where $a = r_1$ and $b = r_2 r_1$. Figure 1.6 shows the graph $\Gamma' = \Gamma(D_3, \{a, b\})$. The dotted edges without arrows correspond to those edges colored by the generator $a = r_1$ of order 2. The solid edges with arrows correspond to those edges colored by the generator $b = r_2 r_1$ of order 3.

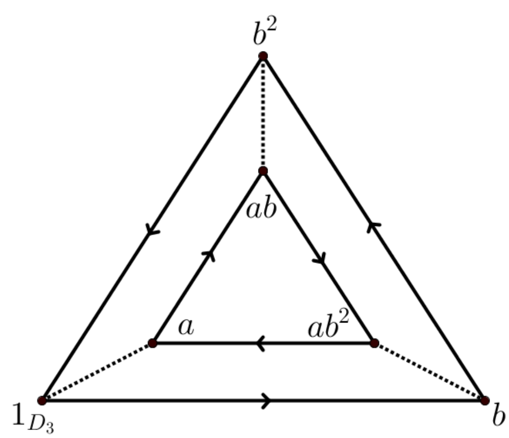

Figure 1.6: Cayley graph $\Gamma(D_3, \{a, b\})$

The left-translation automorphism L_a is not a graph reflection on Γ'. It does have the property that $a^2 = \mathrm{id}_{\Gamma'}$, but the only edge which is inverted by L_a is the edge between 1_{D_3} and a. The removal of this single edge does not separate Γ'. However, if we define a (group) automorphism $\phi : D_3 \to D_3$ by stating that $\phi(a) = a$ and $\phi(b) = b^{-1} = b^2$, then $L_a\phi$ is a graph reflection on Γ'. Specifically, it interchanges the vertices 1_{D_3} and a, ab^2 and b, and ab and b^2. The edges between the pairs of vertices just listed are exactly those edges in Γ' which are inverted by $L_a\phi$. These correspond to the dotted edges in Figure 1.6. The removal of these edges clearly separates the graph Γ'. Notice that the action of $L_a\phi$ on Γ' reverses the direction of the arrows on the edges colored by b. For example, the arrow on the edge between 1_{D_3} and b^2 in the graph Γ' points toward 1_{D_3}, but the arrow on the edge between $L_a\phi(1_{D_3}) = a$ and $L_a\phi(b^2) = ab$ points away from $L_a\phi(1_{D_3})$.

The reflection L_{r_1} on the graph $\Gamma = \Gamma(D_3, \{r_1, r_2\})$ sent solid edges to solid edges and dashed edges to dashed edges. Furthermore, it did not change the direction of the arrows on the edges. (This is a rather trivial observation in this case, since the edges in Γ did not have arrows.) We will say that L_{r_1} is *color fixing*. The graph reflection $L_a\phi$ on $\Gamma' = \Gamma(D_3, \{a, b\})$ sent dotted edges to dotted edges and solid edges to solid edges. However, it reversed the direction arrows on the solid edges. We will say that $L_a\phi$ is *color preserving*. In general, every color-fixing graph automorphism is also a color-preserving graph automorphism.

In the case of the dihedral group D_3 with the generating set $\mathcal{T} = \{a, b\}$, $L_a\phi$ as just defined in the only color-preserving graph reflection on $\Gamma' = \Gamma(D_3, \mathcal{T})$. In fact, it is the only graph reflection of any kind acting on Γ'.

The a-colored edges are all inverted by $L_a\phi$, but the b-colored edges are never inverted by any color-preserving graph reflection acting on Γ'. For example, if ϕ was a color-preserving graph automorphism on Γ' which inverted the edge between 1_{D_3} and b, then ϕ could not invert any of the other edges incident to the 1_{D_3} or b vertices. It might be possible for ϕ to invert one of the edges in the inner triangle, but then it could not invert the other edges in the inner triangle. Removing just one edge from each of the inner and outer triangles will not separate the graph. Therefore, ϕ could not be a graph reflection.

Generalizing from the specific argument just given, we can see that there is no color-preserving graph reflection acting on Γ' which inverts a b-colored edge. Therefore, we will say that the pair, or *group system*, $(D_3, \mathcal{T})$ is not *totally reflected*. However, the group system $(D_3, \mathcal{S})$, where $\mathcal{S} = \{r_1, r_2\}$, is *totally reflected*. Every edge in the graph $\Gamma = \Gamma(D_3, \mathcal{S})$ is inverted by one of the color-fixing, and therefore color-preserving, reflections L_{r_1}, L_{r_2}, or L_{r_3} acting on Γ.

D_3 is just one example of a finite dihedral group. For any positive integer $n \geq 3$, we know that the dihedral group D_n can be generated by two elements r_1 and r_2 corresponding to reflections in $\mathbb{R}^2$ across lines l_1 and l_2, respectively, where the lines l_1 and l_2 meet at the origin and have a pair of vertical angles with measure π/n. As in the case of D_3, the group system $(D_n, \mathcal{S})$, where $\mathcal{S} = \{r_1, r_2\}$, is totally reflected. L_{r_1} and L_{r_2} are color-preserving graph reflections on $\Gamma(D_n, \mathcal{S})$, and together they generate all possible color-preserving graph reflections on $\Gamma = \Gamma(D_n, \mathcal{S})$. For any edge e in the graph Γ, one of the aforementioned color-preserving reflections will invert the edge e.

The dihedral group D_n can also be defined in terms of the following group presentation:

$$D_n = \langle r_1, r_2 \mid (r_1)^2 = (r_2)^2 = (r_1 r_2)^n = 1_{D_n} \rangle$$

Some readers will recognize this as a Coxeter presentation. Therefore, each dihedral group D_n for $n \geq 3$ is a Coxeter group. The pair $(D_n, \mathcal{S})$, where $\mathcal{S} = \{r_1, r_2\}$, is an example, then, of a Coxeter system.

The finite dihedral groups make up just one family of finite Coxeter groups. The renowned geometer H.S.M. Coxeter himself classified the finite Coxeter groups in 1935 [2]. There are also infinite Coxeter groups. One might wonder: Are general Coxeter systems totally reflected? The answer is a resounding "yes." In fact, Coxeter systems serve as our prototype of totally reflected systems. In 1934, H.S.M. Coxeter showed that there was a strong connection between Euclidean reflection groups and Coxeter groups [1]. For this reason, the elements of the generating set $\mathcal{S}$ in a Coxeter system $(W, \mathcal{S})$ are sometimes referred to as *fundamental reflections*. The use of the word *reflection* here is in the Euclidean sense, but it is true that the elements of $\mathcal{S}$ also act by left translation as color-fixing reflections on the Cayley graph $\Gamma(W, \mathcal{S})$ as well [9]. For any $s \in \mathcal{S}$, one can easily observe that the edge between the elements 1_W and s in $\Gamma(W, \mathcal{S})$ is inverted by L_s. In Chapter 5, we will show that a group system is totally reflected as long as for each edge e with initial vertex 1_G there is a color-preserving graph reflection inverting e. Therefore, based on our preceding observations, $(W, \mathcal{S})$ is a totally reflected system.

In the case of the Coxeter system $(W, \mathcal{S})$, every color-preserving graph reflection on $\Gamma(W, \mathcal{S})$ is color fixing. However, our definition of *totally reflected* is based on the slightly weaker *color preserving* graph reflections. There are interesting group systems $(G, \mathcal{S})$ such that the Cayley graph $\Gamma(G, \mathcal{S})$ has no color-fixing reflections but has many color-preserving reflections.

For example, consider the group system $(G, \mathcal{S})$ where $G = \mathbb{Z} = \langle a \rangle$ and $\mathcal{S} = \{a\}$. The Cayley graph $\Gamma(G, \mathcal{S})$ is shown in Figure 4.1. For any $g \in \mathbb{Z}$, $g \neq 1_{\mathbb{Z}}$, the left translation L_g translates the entire Cayley graph $\Gamma(G, \mathcal{S})$ n "edge lengths" to the right or left for some nonzero integer n. The automorphism L_g on $\Gamma(G, \mathcal{S})$ does not invert any edges in the graph and thus cannot possibly be a reflection. However, there are many color-preserving graph reflections acting on $\Gamma(G, \mathcal{S})$. In fact, for any edge e in the graph, there is some color-preserving reflection on $\Gamma(G, \mathcal{S})$ which inverts e. Therefore, $(G, \mathcal{S}) = (\mathbb{Z}, \{a\})$ is totally reflected.

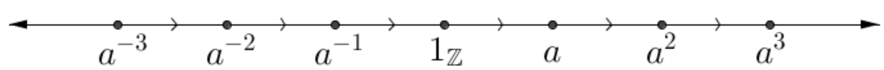

Figure 1.7: Cayley graph $\Gamma(\mathbb{Z}, \{a\})$

By defining the condition of *totally reflected* in terms of color-preserving, and not color-fixing, graph reflections, we have opened ourselves up to a much deeper and more interesting exploration. Indeed, our primary goal in this exposition is to examine the property of *totally reflected*, to develop some basic theory that will expand our understanding of this property, and to try to find larger and larger classes of groups which have this property. In many instances, simple examples will chart a course toward general principles and more complex examples. We will discover ways to combine totally reflected groups to form larger and more interesting totally reflected groups. Finally, we will examine a property, which we will call *strongly totally reflected (or s.t.r.)*, which intuitively seems to be much stronger than the totally reflected (or t.r.) property. We will give examples of groups which are s.t.r., and we will explore the relationship between t.r. groups and s.t.r. groups. In particular, we will try to understand sufficient conditions for a t.r. group to be s.t.r.

In Chapter 2, we will give careful definitions of the terms *graph* and *Cayley graph*, and we will introduce important terminology related to graphs that we will use throughout this and later chapters. We will give examples of Cayley graphs for specific groups and will discuss the general conventions that we will use when drawing Cayley graphs.

In Chapter 3, we will define the concepts of *graph automorphisms* and *groups acting on graphs*. We will then study special types of automorphisms on Cayley graphs, namely, *color fixing* and *color preserving*. We will prove that both color-fixing and color-preserving automorphisms on a Cayley graph $\Gamma(G, \mathcal{S})$ have nice characterizations in terms of simpler types of graph automorphisms.

In Chapter 4, we will delve into the theory concerning graph reflections. We will briefly discuss the reflection groups and their connection to Coxeter groups.

In Chapter 5, we begin our investigation of *totally reflected groups*. First, we will define what it means for a *group system* to be totally reflected. Later, we will define what it means for a group to be totally reflected. Along the way we will give several fundamental examples of totally reflected group systems. We will then begin our investigation into methods of combining totally reflected groups to form larger classes of totally reflected groups. Specifically, we will show that any finite direct or free product of t.r. groups is again t.r. We will also generalize from some of our basic examples involving finite and infinite cyclic groups to describe exactly what conditions must be satisfied in order for a finitely generated abelian group to be t.r.

In Chapter 6, we will generalize the direct and free product constructions to form what we will call a *right-angled product*. We will define a *universal mapping property* for right-angled products. This universal mapping property will make it easier to define right-angled products and to recognize certain groups as right-angled products. We will then state and prove our paramount result, which shows that the right-angled product of totally reflected groups is totally reflected.

In Chapter 7, we will define the concept of *strongly totally reflected (or s.t.r.)*. First, we will define what it means for a group system to be s.t.r., and then we will define what it means for a group to be s.t.r. We will give several examples of s.t.r. groups. We will prove a proposition which will make it easier to show that a reflection group acting on a Cayley graph $\Gamma(G, \mathcal{S})$ has the necessary property which allows us to say that $(G, \mathcal{S})$ is s.t.r. We then turn our attention to the problem of finding sufficient conditions for a t.r. group to be s.t.r. We will use some of these results to show that right-angled Artin groups are s.t.r. This will allow us to define a color-preserving reflection group A_r acting on the Cayley graph of a right-angled Artin group system $(A_\Delta, \mathcal{S})$ such that the chambers of the action of A_r on $\Gamma = \Gamma(A_\Delta, \mathcal{S})$ are the vertices of Γ. We will observe that A_r is isomorphic to a right-angled Coxeter group and will give a complete description in terms of a group presentation for this Coxeter group. We will finish our discussion of right-angled Artin groups by showing that the groups A_Δ and A_r are commensurable. In particular, we will prove that both A_Δ and A_r are finite-index subgroups of the group of color-preserving graph automorphisms of $\Gamma(A_\Delta, \mathcal{S})$. We close the chapter with an example illustrating how the right-angled product graph Δ' for A_r can be derived from the right-angled product graph Δ for A_Δ.

Before we begin, let us quickly make some remarks regarding notation. For any integer n, we will use the symbol I_n to denote the set of the first n positive integers. That is, $I_n := \{i \in \mathbb{Z} \mid 1 \leq i \leq n\}$. Occasionally we will want to include 0 as a possible index. In this case, we will use the symbol $\overline{I}_n$ to denote the set of the first $n+1$ nonnegative integers, starting with 0. That is, $\overline{I}_n = I_n \cup \{0\}$. When working with a group G, we will typically use the symbol 1_G to denote the (unique) identity element for the group. For any element g in the group G, we will use the symbol g^{-1} to denote the (unique) inverse element for g. In any group, the identity element 1_G is its own inverse. For the sake of convenience, we will often write the product $g \cdot h$ of group elements simply as gh. Throughout this exposition, we will make use of standard notation and terminology from group theory. Any conventions we use which may not be standard to the field will be elucidated.

Chapter 2

Graphs

2.1 Definitions

A **graph**, Γ, is a quadruple $\Gamma = (V, E, -, \iota)$, where $V = V(\Gamma)$ and $E = E(\Gamma)$ are sets; $- : E \to E$ is an involution such that for any $e \in E$, $\overline{e} \neq e$; and $\iota : E \to V$ is a function. V is called the **vertex set** of Γ and consists of elements called **vertices** of Γ. E is called the **edge set** of Γ and consists of elements called **edges** of Γ. For any edge e in Γ, we say that the edge $\overline{e}$ is the **reverse** of e, and we will call the function $- : E \to E$ the **reversing function**. The function $\iota : E \to V$ is called the **initial-vertex function**. From this definition of graph, we may derive another function, namely the **terminal-vertex function**, $\tau : E \to V$ defined by $\tau(e) = \iota(\overline{e})$. A **loop (or loop edge)** is an edge e with $\iota(e) = \tau(e)$.

A **subgraph** Γ' of a graph $\Gamma = (V, E, \iota, -)$ is a graph $\Gamma' = (V', E', \iota', -')$, where $V' \subseteq V$ and $E' \subseteq E$, with $\iota(e)$, $\tau(e) \in V'$ and $\overline{e} \in E'$ for all $e \in E'$, and where $\iota' = \iota|_{E'}$ and $-' = -|_{E'}$. In this situation, we write $\Gamma' \subseteq \Gamma$.

We will say that a graph is **simplicial** if no two distinct edges share both initial and terminal points. In more formal terms, a graph is simplicial if whenever $\iota(e) = \iota(e')$ and $\tau(e) = \tau(e')$ for $e, e' \in E$, then $e = e'$. By virtue of the definition, a simplicial graph contains no loop edges. Graph theorists may use the term *simple* in place of *simplicial* in this context. However, we choose to use the more strongly topological term *simplicial* here to reflect the connection between a simplicial graph and a simplicial complex. Indeed, any simplicial graph gives rise to a 1-dimensional CW-complex.

We may define a graph in an equivalent way, by saying that a graph consists of a set of vertices $V = V(\Gamma)$ and a multiset of edges $E = E(\Gamma)$ whose elements are ordered pairs of distinct vertices such that (u, v) is in E if and only if (v, u) is in E. This definition can be seen to be consistent with our previous definition if for any edge (u, v) in the graph Γ (as given by the current definition), we define $\overline{(u, v)} = (v, u)$ and $\iota\big((u, v)\big) = u$. From this perspective, we can say that a graph is simplicial if the edge multiset is specifically a set.

A **(nonempty) path**, say γ, in a graph Γ is a sequence of one or more edges $(e_1, e_2, \ldots, e_n)$, such that whenever $n > 1$ we have $\tau(e_{i-1}) = \iota(e_i)$ for all $i \in I_n - \{1\}$. The nonnegative integer n is called the **length** (or **edge length**) of path γ. For any (nonempty) path γ, the vertex $\iota(e_1)$ is the **initial vertex of the path**, which we will denote by $\iota(\gamma)$. The vertex $\tau(e_n)$ is the **terminal vertex of the path**, which we will denote as $\tau(\gamma)$. We say that the path **passes through** the vertices $\iota(e_1), \iota(e_2), ..., \iota(e_n), \tau(e_n)$. In some instances, it will be convenient for us to consider a single vertex v_0 in Γ to be a path of length $n = 0$. Such a path will be called the **empty path based at v_0**. The vertex v_0 will be the both the initial and terminal points of this empty path.

Given a path $\gamma = (e_1, \ldots, e_n)$, a **subpath** of γ is any path formed by removing the first i and the last j edges of γ, for some nonnegative integers i and j with $0 \le i + j \le n$. If $i + j = n$, then the subpath formed is the empty path based at one of the vertices through which γ passes. If $i + j = 0$, then the subpath formed is just γ itself.

While the functions ι and τ have only formally been defined on the edge set of Γ, it should now be clear how we may extend these functions to paths in Γ. Similarly, we may extend the function $-$ to paths, by defining the **reverse of path $\boldsymbol{\gamma}$** to be the path $\overline{\gamma}$ such that

$$\overline{\gamma} = \begin{cases} \gamma & \text{if } n = 0 \\ (\overline{e_n}, \overline{e_{n-1}}, \ldots, \overline{e_1}) & \text{if } n > 0 \end{cases}$$

Note that $\overline{\gamma}$ is in fact a path. When $n = 0$, this is clear from the definition of $\overline{\gamma}$. When $n = 1$, $\overline{\gamma} = \overline{e_1}$ which is a path. When $n > 1$, $\overline{\gamma}$ is a path since $\tau(\overline{e_i}) = \iota(e_i) = \tau(e_{i-1}) = \iota(\overline{e_{i-1}})$, for all $i \in I_n - \{1\}$. The initial vertex of $\overline{\gamma}$ is $\iota(\overline{\gamma}) = v_0$ if $n = 0$ and is $\iota(\overline{\gamma}) = \iota(\overline{e_n})$ if $n > 0$. More generally, we can see that whether $n = 0$ or $n > 0$, we have $\iota(\overline{\gamma}) = \tau(\gamma)$ and $\tau(\overline{\gamma}) = \iota(\gamma)$.

Whenever γ and δ are paths in a graph Γ with $\tau(\gamma) = \iota(\delta)$, there is a natural way to combine γ and δ into one path from $\iota(\gamma)$ to $\tau(\delta)$. To make this more precise, let $\gamma = (e_1, \ldots, e_n)$ and $\delta = (\widetilde{e_1}, \ldots, \widetilde{e_m})$ be nonempty paths in Γ with $\tau(\gamma) = \iota(\delta)$. We then define the **concatenation** of γ and δ to be the path $\gamma \cdot \delta = (e_1, \ldots, e_n, \widetilde{e_1}, \ldots, \widetilde{e_m})$, which is a well-defined path since $\tau(e_n) = \tau(\gamma) = \iota(\delta) = \iota(\widetilde{e_1})$. If γ is an empty path based at $\iota(\delta)$, then $\gamma \cdot \delta = \delta$. If δ is an empty path based at $\tau(\gamma)$, then $\gamma \cdot \delta = \gamma$.

A path γ is said to be **closed** whenever $\iota(\gamma) = \tau(\gamma)$. A **cycle** (of **length**

$\boldsymbol{n}$) is a closed path $\gamma = (e_1, e_2, ..., e_n)$ satisfying the additional condition that whenever $n \geq 2$ and $i \neq j$, $\iota(e_i) \neq \iota(e_j)$. According to this definition, a loop edge can be thought of as a cycle of length 1, though this type of cycle is not very interesting. It should be noted that a cycle in a simplicial graph must have at least three edges. A cycle of length 3 will be called a **triangle**.

For vertices v and w in Γ, we say that v and w are **adjacent** if there exists a non-loop edge $e \in E(\Gamma)$ with $\iota(e) = v$ and $\tau(e) = w$. In this instance, we say that v and w are **connected by the edge $\boldsymbol{e}$**. More generally, we could talk about vertices which are connected via a path, where that path may or may not be a single edge. For vertices v and w in Γ, a **path from $\boldsymbol{v}$ to $\boldsymbol{w}$** is any path γ with $\iota(\gamma) = v$ and $\tau(\gamma) = w$. In this instance, we say that w is **reachable** from v, or that v and w are **connected by the path $\boldsymbol{\gamma}$**. Any vertex v is reachable from itself by way of the empty path based at v. If w is reachable from v by way of a path γ, then v is reachable from w by way of the path $\overline{\gamma}$. If v_2 is reachable from v_1 and if v_3 is reachable from v_2, then by concatenating paths we can see that v_3 is reachable from v_1. Together, these observations show that reachability is, in fact, an equivalence relation. The equivalence classes under this equivalence relation are called the **connected components of $\boldsymbol{\Gamma}$**. Said another way, the connected components of Γ are subgraphs Γ_1, Γ_2, Γ_3,..., such that each Γ_i is connected and such that no path exists in Γ between v and w for any $v \in V(\Gamma_i)$, $w \in V(\Gamma_j)$ for $i \neq j$.

We say that Γ is **connected** if it has only one connected component. If Γ has more than one connected component, we say that it is **disconnected**. Said more simply, Γ is connected if and only if for any vertices v and w in Γ, there exists a path γ in Γ from v to w. In subsequent chapters, we will consider the effect that removing certain edges will have on the connectedness of Γ. If $\Gamma = (V, E, -, \iota)$ is connected, a subset $E' \subseteq E$ is said to **separate** Γ if the new graph $\Gamma' = (V, E - E', -, \iota)$ is disconnected.

2.2 Cayley Graphs

Groups often are studied as algebraic objects. However, groups can be examined from a more geometric point of view. We will begin by describing a way to represent groups graphically. Consider a group G and a subset $\mathcal{S} \subseteq G$. Let $\mathcal{S}^{\pm} := \mathcal{S} \cup \mathcal{S}^{-1}$, where $\mathcal{S}^{-1} = \left\{ s^{-1} \mid s \in \mathcal{S} \right\}$. We define the **Cayley graph of G (with respect to $\mathcal{S}$)** to be a graph $\Gamma = \Gamma(G, \mathcal{S})$ with vertex set $V(\Gamma) = G$, edge set $E(\Gamma) = \left\{ e(g,s) \mid g \in G, s \in \mathcal{S}^{\pm} \right\}$, initial-vertex function defined by $\iota\big(e(g,s)\big) = g$ for any $e(g,s) \in E$, and reversing function defined by $\overline{e(g,s)} = e(gs, s^{-1})$ for any $e(g,s) \in G$.

In order for $\Gamma(G, \mathcal{S})$ to be a graph, according to our definition, we must insist that the identity element of the group not be in the set $\mathcal{S}$. If 1_G was in $\mathcal{S}$, then $e(g, 1_G) \in E(\Gamma)$ for all $g \in G$, and for any such edge we have $\overline{e(g,1_G)} = e(g \cdot 1_G, 1_G^{-1}) = e(g, 1_G)$. However, in our definition of graph, we specified that for every $e \in E(\Gamma)$ we must have $\overline{e} \neq e$. Therefore, we will assume that $1_G \notin \mathcal{S}$ whenever we are working with a Cayley graph $\Gamma(G, \mathcal{S})$. If the group G is itself the trivial group, meaning that G consists of just one element and which we typically denote as $G = \{1_G\}$, then the assumption we just made that $1_G \notin \mathcal{S}$ could leave us in a strange situation of having the set $\mathcal{S} = \varnothing$. While $\mathcal{S} = \varnothing$ presents no problem for the definition of Cayley graph, it does lead to an uninteresting scenario. Indeed, the Cayley graph $\Gamma(\{1_G\}, \varnothing)$ is the graph consisting of one vertex and no edges. Mostly, the case where G

is the trivial group is not a case in which we are interested. Therefore, unless otherwise noted, we will assume that G is a nontrivial group when working with Cayley graphs.

Lemma 2.1. $\Gamma(G,S)$ *is a simplicial graph.*

Proof. Let $e(g,s)$ and $e(g',s')$ be two edges in $\Gamma(G,S)$. Then $g,g' \in G$ and $s,s' \in S^{\pm}$. Suppose that $\iota\big(e(g,s)\big) = \iota\big(e(g',s')\big)$ and $\tau\big(e(g,s)\big) = \tau\big(e(g',s')\big)$. This implies that $g = g'$ and $gs = g's'$ Combining these equalities, we can conclude that $s = s'$. Thus, $e(g,s) = e(g',s')$. Therefore, in $\Gamma(G,S)$, any two edges which share both initial and terminal points are in fact the same edge, meaning that $\Gamma(G,S)$ is a simplicial graph.

□

Whenever S is a nonempty subset of a group G, we use the notation $\langle S\rangle$ to denote the smallest subgroup of G which contains every element of S. We say that $\langle S\rangle$ is the **subgroup generated by S**. This subgroup must contain all finite products of elements of $S^{\pm}$. We will call any such finite product a **word in $S^{\pm}$**. To be more exact, a word in $S^{\pm}$ is any product of the form $s_1 s_2 \dots s_n$, where n is a positive integer and $s_i \in S^{\pm}$ for $i \in I_n$.

Lemma 2.2. *The Cayley graph* $\Gamma(G,S)$ *is connected if and only if* $G = \langle S\rangle$.

Proof. ($\Longrightarrow$) Assume that $\Gamma(G,S)$ is a connected graph. We must show that $G = \langle S\rangle$. Since we already know that $\langle S\rangle \subseteq G$, we must show that $G \subseteq \langle S\rangle$. We have previously established that S must be a subset of G which is nonempty and which does not contain the identity element. Therefore, there exists a non-identity element $s \in S$. The product $s \cdot s^{-1} = 1_G$ must then be in the subgroup $\langle S\rangle$. Suppose now that g is any non-identity element of

G. Since $\Gamma(G,S)$ is a connected graph, there must be a path, γ, between the vertex 1_G and the vertex g in the graph. Since $g \neq 1_G$, γ is a nonempty path. Therefore, there exists a positive integer n, elements $g_1, g_2, \ldots, g_n \in G$ (where specifically $g_1 = 1_G$), and elements $s_1, s_2, \ldots, s_n \in S^{\pm}$ such that $\gamma = (e(g_1, s_1), e(g_2, s_2), \ldots e(g_n, s_n))$. Since γ is a path terminating at the vertex g, we must have that $g = \tau(\gamma) = g_1 s_1 s_2 \ldots s_n = s_1 \ldots s_n$. Therefore, g can be written as a finite product of elements of $S^{\pm}$, meaning that $g \in \langle S \rangle$. Thus, we've now shown that $g \in \langle S \rangle$ for any $g \in G$, meaning that $G = \langle S \rangle$.

($\Longleftarrow$) Assume that $G = \langle S \rangle$. We must show that the graph $\Gamma(G,S)$ is connected. Let $g, h \in G$ be vertices in $\Gamma(G,S)$. If $g = h$, then the empty path based at g is a path from g to h, trivially. So suppose that $g \neq h$. Since S generates G, there exists a positive integer n and elements $s_1, s_2, \ldots s_n$ of $S^{\pm}$ such that $s_1 s_2 \ldots s_n = g^{-1}h$, which implies that $g s_1 s_2 \ldots s_n = h$. Then $\gamma = \big(e(g, s_1), e(gs_1, s_2), e(gs_1 s_2, s_3), \ldots, e(gs_1 s_2 \ldots s_{n-1}, s_n)\big)$ is a path in $\Gamma(G,S)$ with $\iota(\gamma) = g$ and $\tau(\gamma) = g s_1 s_2 \ldots s_n = h$. Thus, we've now shown that there exists a path in $\Gamma(G,S)$ between any two vertices g and h, meaning that $\Gamma(G,S)$ is a connected graph.

□

The preceding lemma gives one example of the kinds of correspondences which exist between the algebraic and geometric properties of a group. The proof of the lemma highlights the correspondence between paths in the Cayley graph $\Gamma(G,S)$ and words in the set $S^{\pm}$. Indeed, given a non-identity element $g \in G$, each path in $\Gamma(G,S)$ between the vertex 1_G and the vertex g will give us a way of writing g as a word in $S^{\pm}$. More generally, given any elements $g, h \in G$ and any path γ in $\Gamma(G,S)$ between the vertices g and h, traversing the edges of γ from g to h will give us a way of writing $g^{-1}h$ as a word in $S^{\pm}$, say $w = w(\gamma)$. In this scenario, we will say that $w(\gamma)$ is the **word (in $\mathcal{S}^{\pm}$) written by path γ**, or in other words, is the **word (in $\mathcal{S}^{\pm}$) corresponding to path γ**.

As a result of our definition for Cayley graph, we notice that for a group G and a subset $\mathcal{S} \subseteq G$, we have that $\Gamma(G, \mathcal{S}) = \Gamma(G, \mathcal{S}^{\pm})$. In fact, if $s, s^{-1} \in \mathcal{S}$ and if $s^{-1} \neq s$, then removing s^{-1} from the set $\mathcal{S}$ does not change the Cayley graph. That is, $\Gamma(G, \mathcal{S}) = \Gamma(G, \mathcal{S} - \{s^{-1}\})$. Similarly, without changing the Cayley graph, we can choose our set $\mathcal{S}$ such that $\mathcal{S} \cap \mathcal{S}^{-1} = \{s \in \mathcal{S} \mid s = s^{-1}\}$. For simplicity, we will say that a subset $\mathcal{S}$ of a group G is **nice** if $1_G \notin \mathcal{S}$ and if $\mathcal{S} \cap \mathcal{S}^{-1} = \{s \in \mathcal{S} \mid s = s^{-1}\}$. Given any nonempty subset $\mathcal{S}$ of a group G, we can create a nice subset $\mathcal{S}'$ by removing from $\mathcal{S}$, if necessary, the identity element and one element of each pair $\{s, s^{-1}\}$, if both s and s^{-1} are in $\mathcal{S}$ and $s \neq s^{-1}$. One can verify that $\langle \mathcal{S} \rangle = \langle \mathcal{S}' \rangle$. Therefore, without loss of generality we may (and will) assume that a generating set $\mathcal{S}$ for a group G is a nice generating set.

2.3 Drawing Graphs

When working with graphs, it is often convenient to visualize a graph pictorially rather than to work abstractly with the graph's set and function descriptions. Consider a graph $\Gamma = (V, E, -, \iota)$. When we draw Γ, we will use dots to represent the vertices of Γ (with the dots and vertices corresponding bijectively) and we will use curves to represent the edges of Γ (with the edges and curves corresponding bijectively). We will usually label each vertex (dot) with its corresponding element of $V(\Gamma)$. Similarly, we can label each edge (curve) with its corresponding element of $E(\Gamma)$, though in practice we often skip labeling the edges in this way. For any edge e, we place an arrow on its corresponding curve which points from $\iota(e)$ to $\tau(e)$, thereby endowing that edge with a direction.

Most of the graphs we draw will correspond to Cayley graphs of the form $\Gamma = \Gamma(G, \mathcal{S})$, where G is a nontrivial group and $\mathcal{S}$ is a nice generating set. For the sake of simplicity and utility, there are certain conventions we will use when drawing Cayley graphs. Since $V(\Gamma) = G$, the graph will consist of $|G|$ vertices, which we will label with the group elements.

One of the conventions we will use when drawing a Cayley graph depends on what we will call a *coloring* of the edges. Let us first define the color of an edge in $\Gamma(G, \mathcal{S})$. Recall that every edge in $\Gamma(G, \mathcal{S})$ is of the form $e = e(g, s)$, where $g \in G$ and $s \in \mathcal{S}^{\pm}$. In this setting, we say that s is the **color** of the edge $e = e(g, s)$. A **coloring** for $\Gamma(G, \mathcal{S})$ is a subset E' of $E = E(\Gamma)$ which contains exactly one element of each edge pair $\{e, \overline{e}\}$. The coloring that we most often will choose is $E' = \{e(g, s) \mid s \in \mathcal{S}\}$. This is a well-defined coloring since we are assuming that $\mathcal{S}$ is a *nice* generating set. We will call this the **standard coloring** for $\Gamma(G, \mathcal{S})$.

When drawing a Cayley graph $\Gamma(G, \mathcal{S})$, we will only include the directed curves corresponding to the edges in the standard coloring E' as described above. We will label each of these edges with its corresponding color. Sometimes such a label will come in the form of a symbol placed above or near the corresponding curve in $\Gamma(G, \mathcal{S})$. More often, though, such a label will be encoded in the graph visually according to some prescribed legend. For example, we may choose to use unbroken curves to represent all edges colored by the generator s and instead use dotted curves to represent all edges colored by a different generator t.

The convention of drawing and labeling edges as just described allows us to convey important group information in a clear and concise way. For example, in the drawing of $\Gamma(G, \mathcal{S})$, suppose that we start at some vertex g and move along an s-colored edge, in the direction indicated by the arrow, until we reach another vertex h. Then in terms of the group elements and operation, we can say that $gs = h$. But suppose instead that in moving along the s-colored edge from g to h we had to travel in the opposite direction as that indicated with the arrow. Then we can say that $gs^{-1} = h$.

The fact that moving along an edge results in a right multiplication by a generator (or the inverse of a generator, depending on the direction we move with respect to the arrow on the edge) is consistent with our previous definitions involving Cayley graphs and with our notion of path concatenation. The paths in the pictorial representation of $\Gamma(G, \mathcal{S})$ correspond to words in $\mathcal{S}^{\pm}$ just as described immediately following Lemma 2.2.

One more convention we will use when drawing Cayley graphs relates to our treatment of edges which are colored by generators of order 2. For simplicity, we typically leave off the arrow from an edge when its color is given by a generator s of order 2 in the group, since moving in either direction along the arrow would result in a right multiplication by s, since $s^2 = 1_G$ implies that $s = s^{-1}$.

An example will help to illustrate the conventions we use when drawing Cayley graphs. Consider the dihedral group $G = D_4$ with presentation $\langle a, b \mid a^4 = b^2 = 1_{D_4}, bab = a^{-1} \rangle$ and generating set $\mathcal{S} = \{a, b\}$. Figure 2.1 shows the Cayley graph $\Gamma = \Gamma(G, \mathcal{S})$.

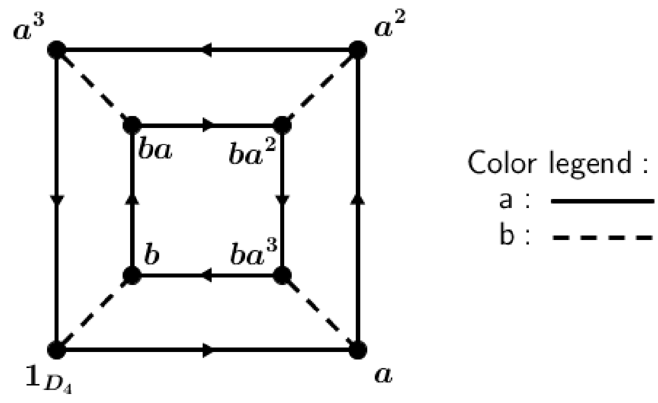

Figure 2.1: Cayley graph $\Gamma(D_4, \{a, b\})$

Paths in the Cayley graph with the same initial and terminal points correspond to words in $\mathcal{S}^{\pm}$ which are the same as group elements. For example, in Figure 2.1, consider two possible paths between the vertex labeled 1_{D_4} and the vertex labeled a^3. The first path, γ_1, begins at the 1_{D_4} vertex, then goes to the b vertex, then to the ba vertex, and finally ends at the a^3 vertex. The word in $\mathcal{S}^{\pm}$ written by path γ_1 is $w(\gamma) = bab$, since from left to right these are the colors of the edges we traversed in going from 1_{D_4} to a^3 and since we traversed each edge in the direction indicated by the arrow. The second path, γ_2, begins at the 1_{D_4} vertex and then goes immediately to the a^3 vertex. The word in $\mathcal{S}^{\pm}$ written by path γ_2 is $w(\gamma_2) = a^{-1}$, since we traversed just a single a-colored edge, but we went in the opposite direction as the arrow drawn. Since these two paths share initial and terminal vertices, we know that their corresponding words must be the same. Therefore, $bab = a^{-1}$, which is, in fact, one of the defining relations for the D_4 group with the given presentation.

A Cayley graph concisely encodes a large quantity of information about the group in question. It helps us to understand the structure of the group in a way that may be difficult to access in a more abstract setting. If $\mathcal{S}$ and $\mathcal{T}$ are different generating sets for a group G, the definition of Cayley graph implies that $\Gamma(G, \mathcal{S})$ and $\Gamma(G, \mathcal{T})$ are different. They have the same vertex set, of course, but have different edge sets.

Consider again the group $G = D_4$ with the same presentation as before: $D_4 = \langle a, b \mid a^4 = b^2 = 1_{D_4}, bab = a^{-1} \rangle$. This time, however, take the generating set to be $\mathcal{T} = \{b, ab\}$. Figure 2.2 shows the Cayley graph $\Gamma(G, \mathcal{T}) = \Gamma(D_4, \{b, ab\})$. For the sake of consistency, the vertices are labeled using the same 8 symbols (in terms of a and b) as used in Figure 2.1. Note that Figure 2.2 also depicts the Cayley graph $\Gamma(D_4, \{c, d\})$, where D_4 is presented as $D_4 = \langle c, d \mid c^2 = d^2 = (cd)^4 = 1_{D_4} \rangle$. One can check that these two presentations do, in fact, give isomorphic groups (which we simply call D_4) by considering the group isomorphism which sends c to ab and which sends d to b.

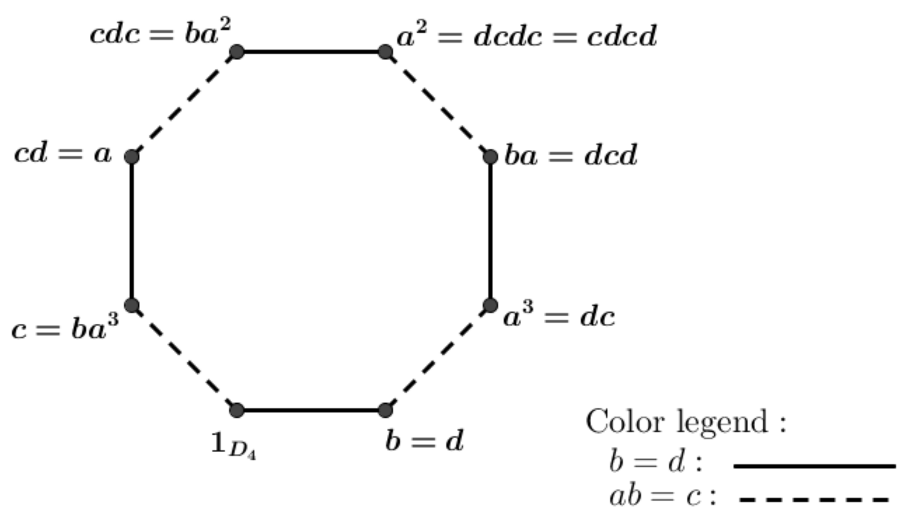

Figure 2.2: Cayley graphs $\Gamma(D_4, \{b, ab\})$ and $\Gamma(D_4, \{c, d\})$

Chapter 3

Graph Morphisms

3.1 Definitions

Suppose now that Γ and Δ are graphs. A **graph morphism** $\phi : \Gamma \to \Delta$ consists of two functions

$$\phi_V : V(\Gamma) \to V(\Delta)$$
$$\phi_E : E(\Gamma) \to E(\Delta)$$

which respect the reversing and initial-vertex functions. More precisely, for all $e \in E(\Gamma)$,

$$\phi_V\big(\iota(e)\big) = \iota\big(\phi_E(e)\big)$$
$$\phi_E(\overline{e}) = \overline{\phi_E(e)}$$

To simplify the notation, when working with a graph morphism $\phi : \Gamma \to \Delta$, we will just write ϕ in place of ϕ_V and ϕ_E when the context is clear. In other words, for $e \in E(\Gamma)$ and $v \in V(\Gamma)$, we'll write $\phi_V(v) = \phi(v)$ and $\phi_E(e) = \phi(e)$. Suppose now that we have another graph morphism, $\psi : \Delta \to \Lambda$. We define the **composition** $\psi \circ \phi : \Gamma \to \Lambda$ to be the graph morphism consisting of the two functions

$$(\psi \circ \phi)_V = \psi_V \circ \phi_V : V(\Gamma) \to V(\Lambda)$$
$$(\psi \circ \phi)_E = \psi_E \circ \phi_E : E(\Gamma) \to E(\Lambda)$$

Usually we will omit the composition symbol and write $\psi\phi$ instead of $\psi \circ \phi$, unless the omission would lead to confusion.

In some settings, we may want to use a slightly weaker type of morphism between graphs, one which allows the "collapsing" of edges to vertices. A **non-rigid graph morphism** $\phi : \Gamma \to \Delta$ consists of two functions

$$\phi_V : V(\Gamma) \to V(\Delta)$$
$$\phi_E : E(\Gamma) \to E(\Delta) \cup V(\Delta)$$

which respect the reversing and initial-vertex functions. More precisely, for all $e \in E(\Gamma)$,

$$\phi_V\big(\iota(e)\big) = \begin{cases} \iota\big(\phi_E(e)\big) & \text{if } \phi_E(e) \in E(\Delta) \\ \phi_E(e) & \text{if } \phi_E(e) \in V(\Delta) \end{cases}$$
$$\phi_E(\overline{e}) = \begin{cases} \overline{\phi_E(e)} & \text{if } \phi_E(e) \in E(\Delta) \\ \phi_E(e) & \text{if } \phi_E(e) \in V(\Delta) \end{cases}$$

As before, we will often just write ϕ in place of both ϕ_V and ϕ_E, when the context is clear. Any graph morphism is also a non-rigid graph morphism in which the image of ϕ_E has empty intersection $V(\Delta)$. For the sake of clarity, we will use the term *graph morphism* to mean a non-rigid graph morphism $\phi : \Gamma \to \Delta$ such that $\phi(e) \in E(\Gamma)$ for all $e \in E(\Gamma)$.

An **embedding** of one graph, Γ, into another graph, Δ, is a graph morphism $\phi : \Gamma \to \Delta$ such that the associated functions ϕ_V and ϕ_E are injective. A **graph isomorphism** between Γ and Δ is a graph morphism $\phi : \Gamma \to \Delta$ such that the associated functions ϕ_V and ϕ_E are both bijective. Given any two graphs Γ and Δ, we say that they are **isomorphic (graphs)** if there exists a graph isomorphism between Γ and Δ. A **graph automorphism** of Γ is a graph isomorphism $\phi : \Gamma \to \Gamma$. In this context, we will sometimes say that ϕ is an automorphism *acting on* (or just *on*) Γ. The set of all graph automorphisms on Γ is denoted $\mathrm{Aut}(\Gamma)$.

Lemma 3.1. $\mathrm{Aut}(\Gamma)$ *is a group under the operation of composition, called the* ***automorphism group of*** $\boldsymbol{\Gamma}$*.*

The proof of this lemma is a standard result in graph theory. The identity element of the group $\mathrm{Aut}(\Gamma)$ is the identity automorphism, which fixes all vertices and edges. We will use the symbol id_Γ to denote the identity automorphism on the graph Γ.

Perhaps a more natural way to define graph automorphism involves permutations. A graph automorphism on Γ can be thought of as a permutation σ on $V(\Gamma)$ which preserves the structure of the graph. This means that if there is an edge in Γ with initial point v and terminal point w, then there is an edge in Γ with initial point $\sigma(v)$ and terminal point $\sigma(w)$. This characterization of graph automorphism is more intuitive and useful than the formal definition given previously.

3.2 Automorphisms of Cayley Graphs

In the preceding section, we defined the concept of graph automorphism. In this section we will examine graph automorphisms specifically on Cayley graphs. Assume that G is a group with a generating set $\mathcal{S}$. Recall that we are assuming that G is nontrivial and $\mathcal{S}$ is a nice generating set. Consider the Cayley graph $\Gamma = \Gamma(G, \mathcal{S})$.

Recall that any edge in Γ must be of the form $e(g,s)$, where $g \in G$ and $s \in \mathcal{S}^{\pm}$. Any graph automorphism ϕ on Γ must send $e(g,s)$ to another edge in Γ, say $e(h,t)$, where $h \in G$ and $t \in \mathcal{S}^{\pm}$. Since $e(g,s)$ has initial point g and terminal point gs, the edge $\phi\big(e(g,s)\big)$ must have initial point $\phi(g)$ and terminal point $\phi(gs)$. Therefore, we must have $h = \phi(g)$ and $ht = \phi(g)t = \phi(gs)$, and so $t = \phi(g)^{-1}\phi(gs)$. Thus, $\phi\big(e(g,s)\big) = e\big(\phi(g), \phi(g)^{-1}\phi(gs)\big)$. In order for $\phi\big(e(g,s)\big)$ to be a well-defined edge, we must have that $\phi(g)^{-1}\phi(gs)$ is an element of $\mathcal{S}^{\pm}$. More broadly, any graph automorphism ϕ on Γ must satisfy the following condition:

$$\phi(g)^{-1}\phi(gs) \in \mathcal{S}^{\pm}, \qquad \text{for any } g \in G, s \in \mathcal{S}^{\pm}$$

This formula is not particularly nice, but if we knew that $\phi_V : G \to G$ was a group automorphism of G, then we would have more simply that $\phi\big(e(g,s)\big) = e\big(\phi(g), \phi(s)\big)$.

In general, we must be careful to remember that not every graph automorphism ϕ on $\Gamma(G, \mathcal{S})$ satisfies the condition that $\phi_V : G \to G$ is a group automorphism of G. With that said, we will soon see that there is some relationship between graph automorphisms on Γ and group automorphisms of G. We will let $\text{Aut}(G)$ denote the set of all automorphisms of the group G. Like $\text{Aut}(\Gamma)$, $\text{Aut}(G)$ takes on its own group structure.

Lemma 3.2. $\text{Aut}(G)$ *is a group under the operation of composition, called the* ***automorphism group of G****.*

The proof of this lemma is a standard result in group theory. The identity element of the group $\text{Aut}(G)$ is the identity automorphism, which fixes every group element. We will use the symbol id_G to denote the identity automorphism of the group G.

If G is generated by $\mathcal{S}$, then G is also generated by $\phi(\mathcal{S}) = \{\phi(s) \mid s \in \mathcal{S}\}$, where $\phi \in \text{Aut}(G)$. Some, but not all, group automorphisms will fix the set $\mathcal{S}^{\pm}$. The set of all such group automorphisms, ϕ, such that $\phi(\mathcal{S}^{\pm}) = \mathcal{S}^{\pm}$ is a subgroup of $\text{Aut}(G)$, which we will denote by $\text{Aut}(G, \mathcal{S}^{\pm})$ and which we will call the **generating-set-preserving automorphism group**.

Another way to think of the subgroup $\text{Aut}(G, \mathcal{S}^{\pm})$ is in terms of permutations. Any group automorphism ϕ of G is a permutation of the group elements which also preserves the structure of the group, meaning that $\phi(gh) = \phi(g)\phi(h)$ for any group elements g and h. Any member of the subgroup $\text{Aut}(G, \mathcal{S}^{\pm})$ is a permutation of $\mathcal{S}^{\pm}$ which extends to an automorphism of the entire group G. In our work, the generating set $\mathcal{S}$ will typically be small. As such, one can efficiently identify the elements of $\text{Aut}(G, \mathcal{S}^{\pm})$ by examining the permutations on $\mathcal{S}^{\pm}$ and testing whether or not they extend to group automorphisms.

Some, but not all, of the elements of $\text{Aut}(G)$ extend to automorphisms of the Cayley graph $\Gamma = \Gamma(G, \mathcal{S})$. Elements of $\text{Aut}(G)$ preserve the structure of the group G, but we may not be able to extend them in a way that preserves the structure of the graph $\Gamma(G, \mathcal{S})$. A graph automorphism is a function that acts on both vertices and edges. Any group automorphism ϕ of G naturally defines a bijective function on the vertices of $\Gamma(G, \mathcal{S})$, since $V(\Gamma) = G$. We hope to extend ϕ so that it acts also on the edges of Γ in such a way as to ensure that ϕ is a graph automorphism of Γ.

To understand this extension process, suppose that $e = e(g, s)$ is an arbitrary edge in Γ, where $g \in G$ and $s \in \mathcal{S}^{\pm}$. The initial and terminal points of e, respectively, are g and gs. Since g and gs are connected by an edge in Γ, we need $\phi(g)$ and $\phi(gs)$ to be connected by an edge in Γ. Since ϕ is a group automorphism, $\phi(gs) = \phi(g)\phi(s)$. In order for there to be an edge in Γ between $\phi(g)$ and $\phi(g)\phi(s)$, we simply need that $\phi(s) \in \mathcal{S}^{\pm}$. Since the edge e was chosen arbitrarily, we must have more broadly that $\phi(\mathcal{S}^{\pm}) = \mathcal{S}^{\pm}$. In other words, we need for ϕ to be an element of the subgroup $\text{Aut}(G, \mathcal{S}^{\pm})$ of $\text{Aut}(G)$. Lemma 3.3 formalizes these observations.

Lemma 3.3. $\text{Aut}(G, \mathcal{S}^{\pm})$ *is a subgroup of* $\text{Aut}(\Gamma)$*, where* $\Gamma = \Gamma(G, \mathcal{S})$*.*

Proof. First, we will show that $\text{Aut}(G, \mathcal{S}^{\pm}) \subseteq \text{Aut}\Gamma$ Let $\phi \in \text{Aut}(G, \mathcal{S}^{\pm})$. Since ϕ is an automorphism on G and since $V(\Gamma) = G$, we can define $\phi_V : V(\Gamma) \to V(\Gamma)$ by $\phi_V := \phi$. We must now use ϕ to define a function $\phi_E : E(\Gamma) \to E(\Gamma)$. Suppose $e \in E(\Gamma)$ is arbitrary. Then there exists $g \in G$ and $s \in \mathcal{S}^{\pm}$ such that $e = e(g, s)$. Define ϕ_E by stating that $\phi_E(e) = \phi_E\big(e(g, s)\big) = e\left(\phi(g), \phi(s)\right)$. Since $\phi\left(\mathcal{S}^{\pm}\right) = \mathcal{S}^{\pm}$, $e\left(\phi(g), \phi(s)\right)$ is a well-defined edge in Γ.

We must show that ϕ respects the reversing and initial-vertex functions inherent to Γ. Suppose $e(g,s) \in E(\Gamma)$ is arbitrary. Observe the following:

$$\phi_V(\iota(e)) = \phi_V(g) = \phi(g) = \iota\Big(e(\phi(g), \phi(s))\Big) = \iota\Big(\phi_E\big(e(g,s)\big)\Big)$$

Thus, ϕ respects the initial-vertex function of Γ.

$$\begin{aligned}\phi_E\Big(\overline{e(g,s)}\Big) &= \phi_E\big(e(gs, s^{-1})\big) = e(\phi(gs), \phi(s^{-1}))\\ &= e(\phi(g)\phi(s), \phi(s)^{-1})\\ &= \overline{e(\phi(g), \phi(s))} = \overline{\phi_E\big(e(g,s)\big)}\end{aligned}$$

Thus, ϕ respects the reversing function as well as the initial-vertex function of Γ, and so ϕ does extend to a graph morphism from Γ to itself.

Now, we must show that ϕ_V and ϕ_E are bijective functions. The function ϕ_V is clearly bijective since $\phi : G \to G$ is bijective. Suppose $e(g,s), e(h,t) \in E(\Gamma)$ with $\phi_E\big(e(g,s)\big) = \phi_E\big(e(h,t)\big)$. This implies that $e(\phi(g), \phi(s)) = e(\phi(h), \phi(t))$. Consequently, $\phi(g) = \phi(h)$ and $\phi(s) = \phi(t)$. Since $\phi : G \to G$ is an automorphism, and thus is bijective, we can conclude that $g = h$ and $s = t$. Therefore, $e(g,s) = e(h,t)$, and so ϕ_E is injective. Let $e(g,s) \in E(\Gamma)$. Then $e(\phi^{-1}(g), \phi^{-1}(s)) \in E(\Gamma)$ and we can observe the following:

$$\phi_E\big(e(\phi^{-1}(g), \phi^{-1}(s))\big) = e\left(\phi\phi^{-1}(g), \phi\phi^{-1}(s)\right) = e(g,s)$$

This shows that ϕ_E is surjective. Therefore, ϕ_E is a bijection.

In summary, we can extend $\phi \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$ to a graph automorphism on $\Gamma = \Gamma(G, \mathcal{S})$, implying that $\mathrm{Aut}(G, \mathcal{S}^{\pm}) \subseteq \mathrm{Aut}(\Gamma)$. But is $\mathrm{Aut}(G, \mathcal{S}^{\pm})$ a subgroup of $\mathrm{Aut}(\Gamma)$? In order for this to be true, we must have that for any

$\phi, \psi \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$, $\phi\psi^{-1} \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$. This is clearly true since $\mathrm{Aut}(G, \mathcal{S}^{\pm})$ is a group in its own right. Thus, $\mathrm{Aut}(G, \mathcal{S}^{\pm})$ is a subgroup of $\mathrm{Aut}(\Gamma)$.

□

3.3 Groups Acting on Graphs

Suppose G is a group and Γ is a graph. A **group action of G on Γ** is a group homomorphism $\phi : G \to \mathrm{Aut}(\Gamma)$ satisfying the following properties:

(i) $\phi(1_G)(v) = v, \ \forall\ v \in V(\Gamma)$

(ii) for any $g, h \in G$, $\phi(g)\left(\phi(h)(v)\right) = \phi(gh)(v), \ \forall\ v \in V(\Gamma)$

We often will write $\phi_g(v)$ in place of $\phi(g)(v)$. Using this alternate notation, properties (i) and (ii) can be rewritten as shown here:

(i) $\phi_{1_G}(v) = v, \ \forall\ v \in V(\Gamma)$

(ii) for any $g, h \in G$, $\phi_g\phi_h(v) = \phi_{gh}(v), \ \forall\ v \in V(\Gamma)$

If such a group action of G on Γ exists, we say that G **acts** on Γ.

Suppose G acts on Γ. The group action is said to be **vertex free** if for any $v \in V(\Gamma)$ and $g \in G$, $\phi_g(v) = v \iff g = 1_g$. Likewise, the group action is said to be **edge free** if for any $e \in E(\Gamma)$ and $g \in G$, $\phi_g(e) = e \iff g = 1_G$. If for any pair of distinct vertices v_1 and v_2 in Γ there exists $g \in G$ such that $\phi_g(v_1) = v_2$, we say that the group action is **vertex transitive**. Moreover, if the g in the preceding definition is unique, we say that the group action is **simply vertex transitive**. If for any pair of distinct edges e_1 and e_2 in Γ there exists $g \in G$ such that $\phi_g(e_1) = e_2$, we say that the group action is **edge**

transitive. Moreover, if the g in the preceding definition is unique, we say that the group action is **simply edge transitive**.

Suppose that G is a group and $\mathcal{S}$ is a nice generating set for G. Let $\Gamma = \Gamma(G, \mathcal{S})$. Fix $g \in G$. Define a function $L_g : \Gamma \to \Gamma$ by stating that $L_g(v) = gv$ for any $v \in V(\Gamma)$ and $L_g\big(e(g', s)\big) = e(gg', s)$ for any $e(g', s) \in E(\Gamma)$. One can easily check that L_g is a graph automorphism on Γ. Now define a function $L : G \to \mathrm{Aut}(\Gamma)$ by stating that $L(g) = L_g$ for any $g \in G$. Notice that for any $g, h \in G$, $L(gh) = L_{gh} = L_g L_h = L(g)L(h)$. Therefore, L is a group homomorphism.

Lemma 3.4. *The group homomorphism L is a group action of G on $\Gamma = \Gamma(G, \mathcal{S})$. Moreover, this group action is simply vertex transitive and simply edge transitive.*

Proof. For any $v \in V(\Gamma) = G$, notice that $L_{1_G}(v) = 1_G v = v$. Therefore, L satisfies condition (i) of a group action. Suppose $g, h \in G$. Then notice that $L_g L_h(v) = ghv = L_{gh} v$ for any $v \in V(\Gamma) = G$. Therefore, L satisfies condition (ii) of a group action.

Given any two distinct vertices u and v in Γ, there exists a unique element $g \in G$, namely $g = vu^{-1}$, such that $L_g(u) = v$. Thus, the group action is simply vertex transitive.

Suppose that $e(h, s)$ and $e(h', s')$ are distinct edges in Γ and that $g \in G$. Observe that $L_g\big(e(h, s)\big) = e(h', s')$ if and only if $e(gh, s) = e(h', s')$. But in order for this to happen, we must have $gh = h'$ and $s = s'$. Therefore, given any two distinct edges $e(h, s)$ and $e(h', s')$ in Γ, there exists a unique element $g \in G$, namely $g = h'h^{-1}$, such that $L_g\big(e(h, s)\big) = e(h', s')$. Thus, the group action is simply edge transitive. □

When a group action of G on $\Gamma(G, \mathcal{S})$ is defined using the group homomorphism L, we say that **G acts on $\Gamma(G, \mathcal{S})$ by left translation**. For convenience, we will often use the symbol G to denote the set $L(G) = \{L_g \mid g \in G\}$ of left translations of G on Γ. The context will make clear whenever G is being used to denote $L(G)$.

3.4 Color-Fixing Graph Automorphisms on $\Gamma(G, \mathcal{S})$

In Section 2.3, we defined the concept of an edge's *color*. Let us now introduce some notation to go along with this concept. Recall that for an edge $e(g, s)$ in the Cayley graph $\Gamma = \Gamma(G, \mathcal{S})$, we say that s is the color of that edge. Define a function $c : E(\Gamma) \to \mathcal{S}^{\pm}$ by stating that $c\big(e(g, s)\big) = s$. Said differently, the function c maps each edge to its color. We will call this function the **edge-color function**. Note that for any edge $e(g, s)$ we have the following:

$$c\left(\overline{e(g,s)}\right) = c\big(e(gs, s^{-1})\big) = s^{-1} = \Big[c\big(e(g,s)\big)\Big]^{-1}$$

A graph automorphism $\phi \in \mathrm{Aut}(\Gamma)$ will be called a **color-fixing graph automorphism** on Γ if for any edge $e \in E(\Gamma)$ we have that $c(e) = c\big(\phi(e)\big)$. We will denote the set of all color-fixing graph automorphisms on Γ as $\mathrm{Aut}_{\text{c.f.}}(\Gamma)$. Notice that for each $g \in G$, the left-translation L_g is a color-fixing graph automorphism. But do all color-fixing graph automorphisms arise as left-translations by an element $g \in G$? Lemma 3.5 says that the answer to this question is "yes."

Lemma 3.5. $G = \mathrm{Aut}_{\text{c.f.}}(\Gamma)$

Proof. As we've already seen, $G = \{L_g \mid g \in G\}$ is a subgroup of $\mathrm{Aut}(\Gamma)$ acting on $\Gamma(G, \mathcal{S})$ by left-translation. Let g be an arbitrary element of G, and suppose $e(h,s) \in E(\Gamma)$. As previously noted, $L_g\big(e(h,s)\big) = e(gh,s)$, and so L_g is a color-fixing graph automorphism. Therefore, $G = \{L_g \mid g \in G\} \subseteq \mathrm{Aut}_{\text{c.f.}}(\Gamma)$. To show the reverse inclusion, let $\phi \in \mathrm{Aut}_{\text{c.f.}}(\Gamma)$. Suppose $e(g,s)$ is an arbitrary edge in Γ, and suppose $\phi\big(e(g,s)\big) = e(h,s)$. Now,

$$L_{gh^{-1}}\phi\big(e(g,s)\big) = L_{gh^{-1}}\big[e(h,s)\big] = e(gh^{-1}h,s) = e(g,s).$$

This shows that $L_{gh^{-1}}\phi$ fixes all of the edges of Γ, which implies that it fixes the entire graph. Thus, $L_{gh^{-1}}\phi = \mathrm{id}_\Gamma$, and so $\phi = \big(L_{gh^{-1}}\big)^{-1} = L_{hg^{-1}} \in G$. Therefore, $\mathrm{Aut}_{\text{c.f.}}(\Gamma) \subseteq G$, and thus $G = \mathrm{Aut}_{\text{c.f.}}(\Gamma)$.

□

Corollary 3.6. $\mathrm{Aut}_{\text{c.f.}}(\Gamma)$ *is a subgroup of* $\mathrm{Aut}(\Gamma)$.

Proof. Recall the group homomorphism $L : G \to \mathrm{Aut}(\Gamma)$ which defines the group action of G on Γ by left translation. Then $L(G)$ is a subgroup of $\mathrm{Aut}(\Gamma)$, and so G can be viewed as a subgroup of $\mathrm{Aut}(\Gamma)$ by identifying G with $L(G)$. Thus, by the preceding lemma we can conclude that $\mathrm{Aut}_{\text{c.f.}}(\Gamma)$ is a subgroup of $\mathrm{Aut}(\Gamma)$.

□

3.5 Color-Preserving Graph Automorphisms on $\Gamma(G, \mathcal{S})$

A **color-preserving graph automorphism** is an element $\phi \in \mathrm{Aut}(\Gamma)$ which satisfies the property that for any two edges $e_1, e_2 \in E(\Gamma)$, $c(e_1) = c(e_2)$ if and only if $c\big(\phi(e_1)\big) = c\big(\phi(e_2)\big)$. We will denote the set of all color-preserving graph automorphisms by $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$.

By definition, if $\phi \in \mathrm{Aut}_{\text{c.p.}}(\Gamma)$ sends one s-colored edge to a t-colored edge, then ϕ sends all (and only) the s-colored edges to t-colored edges. This is the reason we call ϕ *color preserving*. If $\phi \in \mathrm{Aut}_{\text{c.p.}}(\Gamma)$ sends s-colored edges to t-colored edges, then ϕ also sends all (and only) s^{-1}-colored edges to t^{-1}-colored edges. To see this, suppose that e is an edge with $c(e) = s$ and $c\big(\phi(e)\big) = t$ for $s, t \in \mathcal{S}^{\pm}$. Then $c(\overline{e}) = [c(e)]^{-1} = s^{-1}$ and

$$c\big(\phi(\overline{e})\big) = c\left(\overline{\phi(e)}\right) = \Big[c\big(\phi(e)\big)\Big]^{-1} = t^{-1}$$

It is clear by definition that $\mathrm{Aut}_{\text{c.p.}}(\Gamma) \subseteq \mathrm{Aut}(\Gamma)$. The next lemma establishes a stronger relationship.

Lemma 3.7. $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$ *is a subgroup of* $\mathrm{Aut}(\Gamma)$.

Proof. As mentioned, we know that $\mathrm{Aut}_{\text{c.p.}}(\Gamma) \subseteq \mathrm{Aut}(\Gamma)$. Let $\phi \in \mathrm{Aut}_{\text{c.p.}}(\Gamma)$. First, we must show that $\phi^{-1} \in \mathrm{Aut}_{\text{c.p.}}(\Gamma)$. Suppose that $e_1, e_2 \in E(\Gamma)$. Since ϕ is a graph automorphism, ϕ must act bijectively on the edges of Γ. Thus, there exist edges $\widetilde{e_1}, \widetilde{e_2} \in E(\Gamma)$ such that $\phi(\widetilde{e_1}) = e_1$ and $\phi(\widetilde{e_2}) = e_2$, or equivalently, $\widetilde{e_1} = \phi^{-1}(e_1)$ and $\widetilde{e_2} = \phi^{-1}(e_2)$. By making use of the fact that ϕ

is color preserving, we can observe the following:

$$\begin{array}{ccc} c\big(\phi^{-1}(e_1)\big) = c\big(\phi^{-1}(e_2)\big) & \Longleftrightarrow & c(\widetilde{e_1}) = c(\widetilde{e_2}) \\ & & \Updownarrow \\ c(e_1) = c(e_2) & \Longleftrightarrow & c\big(\phi(\widetilde{e_1})\big) = c\big(\phi(\widetilde{e_2})\big) \end{array}$$

This shows that $c(e_1) = c(e_2)$ if and only if $c\big(\phi^{-1}(e_1)\big) = c\big(\phi^{-1}(e_2)\big)$, which implies that $\phi^{-1} \in \mathrm{Aut}_{\text{c.p.}}(\Gamma)$.

Now, suppose that $\phi, \psi \in \mathrm{Aut}_{\text{c.p.}}(\Gamma)$. We must show that $\phi\psi \in \mathrm{Aut}_{\text{c.p.}}(\Gamma)$. To this end, suppose that $e_1, e_2 \in E(\Gamma)$. Then

$$c(e_1) = c(e_2) \quad \Leftrightarrow \quad c\big(\psi(e_1)\big) = c\big(\psi(e_2)\big) \quad \Leftrightarrow \quad c\big(\phi\psi(e_1)\big) = c\big(\phi\psi(e_2)\big)$$

where the first implication is a result of ψ being color-preserving and the second implication is a result of ϕ being color-preserving. This shows that $\phi\psi \in \mathrm{Aut}_{\text{c.p.}}(\Gamma)$.

We have now shown that $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$ is a subset of $\mathrm{Aut}(\Gamma)$ which is closed under the group operation and closed under taking inverses. Therefore, $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$ is a subgroup of $\mathrm{Aut}(\Gamma)$.

□

Lemma 3.5 gives us a nice characterization of the elements of the group $\mathrm{Aut}_{\text{c.f.}}(\Gamma)$. In particular, any element of $\mathrm{Aut}_{\text{c.f.}}(\Gamma)$ is equal to a left-translation of the form L_g, for some $g \in G$. Likewise, it would be convenient to have an alternate, perhaps simpler, characterization of the elements of the group $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$. Indeed, as Theorem 3.8 will show, any element of $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$ can be expressed in the form $L_g\alpha$, where $g \in G$ and $\alpha \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$.

Theorem 3.8. $\mathrm{Aut}_{\text{c.p.}}(\Gamma) \cong G \rtimes \mathrm{Aut}(G, \mathcal{S}^{\pm})$

Proof.

Step 1: To show that $G = \mathrm{Aut}_{\text{c.f.}}(\Gamma)$ is a normal subgroup of $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$.

First, it is clear from definitions that $\mathrm{Aut}_{\text{c.f.}}(\Gamma) \subseteq \mathrm{Aut}_{\text{c.p.}}(\Gamma)$. Then by Lemma 3.5 and Corollary 3.6, $G = \mathrm{Aut}_{\text{c.f.}}(\Gamma)$ is a subgroup of $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$. We must show that $\mathrm{Aut}_{\text{c.f.}}(\Gamma)$ is a *normal* subgroup of $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$. To this end, let $\alpha \in \mathrm{Aut}_{\text{c.f.}}(\Gamma)$ and $\phi \in \mathrm{Aut}_{\text{c.p.}}(\Gamma)$. We need to show that $\phi^{-1}\alpha\phi \in \mathrm{Aut}_{\text{c.f.}}(\Gamma)$. Suppose $e(g,s) \in E(\Gamma)$. Since ϕ is color-preserving, we know that for some $t \in \mathcal{S}^{\pm}$, ϕ sends s-colored edges to t-colored edges and ϕ^{-1} sends t-colored edges to s-colored edges. Then for some $g' \in G$, $\phi\big(e(g,s)\big) = e(g',t)$. Since α is color-fixing, there exists some $g'' \in G$ with $\alpha\big(e(g',t)\big) = e(g'',t)$. Finally, there exists some $g''' \in G$ with $\phi^{-1}\big(e(g'',t)\big) = e(g''',s)$. Thus, $\phi^{-1}\alpha\phi\big(e(g,s)\big) = e(g''',s)$, implying that $\phi^{-1}\alpha\phi \in \mathrm{Aut}_{\text{c.f.}}(\Gamma)$.

Therefore, $G = \mathrm{Aut}_{\text{c.f.}}(\Gamma)$ is a normal subgroup of $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$.

Step 2: To show that $\mathrm{Aut}(G, \mathcal{S}^{\pm})$ is a subgroup of $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$.

We will first show that $\mathrm{Aut}(G, \mathcal{S}^{\pm}) \subseteq \mathrm{Aut}_{\text{c.p.}}(\Gamma)$. Let $\alpha \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$. Recall that if $e(g,s) \in E(\Gamma)$, then $\alpha\big(e(g,s)\big) = e\big(\alpha(g), \alpha(s)\big)$. Since $e(g,s) \in E(\Gamma)$ was arbitrary, we can conclude that for any edge $e \in E(\Gamma)$ we have $c\big(\alpha(e)\big) = \alpha\big(c(e)\big)$. Suppose $e_1, e_2 \in E(\Gamma)$. Then observe the following:

$$c(e_1) = c(e_2) \iff \alpha\big(c(e_1)\big) = \alpha\big(c(e_2)\big) \iff c\big(\alpha(e_1)\big) = c\big(\alpha(e_2)\big)$$

This shows that α is color-preserving. Thus, $\mathrm{Aut}(G, \mathcal{S}^{\pm}) \subseteq \mathrm{Aut}_{\text{c.p.}}(\Gamma)$. This inclusion, together with Lemmas 3.3 and 3.7, gives us that $\mathrm{Aut}(G, \mathcal{S}^{\pm})$ is a subgroup of $\mathrm{Aut}_{\mathrm{c.p.}}(\Gamma)$.

Step 3: To show that $G \cap \mathrm{Aut}(G, \mathcal{S}^{\pm}) = \{\mathrm{id}_\Gamma\}$.

Lemma 3.3 and Corollary 3.6 tell us that G and $\mathrm{Aut}(G, \mathcal{S}^{\pm})$ are subgroups of $\mathrm{Aut}(\Gamma)$. From this point of view, the identity element id_Γ of $\mathrm{Aut}(\Gamma)$ is an element of both G and $\mathrm{Aut}(G, \mathcal{S}^{\pm})$. In G, $L_{1_G} = \mathrm{id}_\Gamma$. Suppose that for some $g \in G$ we have $L_g \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$. Of course $L_g(1_G) = g$, but since $L_g \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$ we must have that $L_g(1_G) = 1_G$. Thus, $g = 1_G$, and so $L_g = L_{1_G} = \mathrm{id}_\Gamma$.

Therefore, $G \cap \mathrm{Aut}(G, \mathcal{S}^{\pm}) = \{\mathrm{id}_\Gamma\}$.

Step 4: To show that $\mathrm{Aut}_{\mathrm{c.p.}}(\Gamma) = G \cdot \mathrm{Aut}(G, \mathcal{S}^{\pm}) = \{L_g \alpha \mid g \in G \text{ and } \alpha \in \mathrm{Aut}(G, \mathcal{S}^{\pm})\}$.

Since $G = \{L_g \mid g \in G\}$ and $\mathrm{Aut}(G, \mathcal{S}^{\pm})$ are both subgroups of $\mathrm{Aut}_{\mathrm{c.p.}}(\Gamma)$, as previously shown, clearly $G \cdot \mathrm{Aut}(G, \mathcal{S}^{\pm}) \subseteq \mathrm{Aut}_{\mathrm{c.p.}}(\Gamma)$. We must show the reverse inclusion. To this end, let $\phi \in \mathrm{Aut}_{\mathrm{c.p.}}(\Gamma)$. Define $\sigma := L_{\phi(1_G)^{-1}}\phi$. Consider an arbitrary $s \in \mathcal{S}^{\pm}$ and an arbitrary s-colored edge, say $e(g, s)$. Since ϕ is color-preserving, there exists some $t \in \mathcal{S}^{\pm}$ such that ϕ sends all s-colored edges to t-colored edges. So $\phi\big(e(g, s)\big) = e(g', t)$ for some $g' \in G$. Then since $L_{\phi(1_G)^{-1}}$ is color-fixing, $L_{\phi(1_G)^{-1}}\big(e(g', t)\big) = e(g'', t)$ for some $g'' \in G$. Thus, $L_{\phi(1_G)^{-1}}\phi\big(e(g, s)\big) = e(g'', t)$. Therefore, $L_{\phi(1_G)^{-1}}\phi$ sends all s-colored edges to t-colored edges. Since s was chosen arbitrarily, we now know that $\sigma = L_{\phi(1_G)^{-1}}\phi$ is color-preserving.

We claim that $\sigma \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$. First, we will show that $\sigma(\mathcal{S}^{\pm}) = \mathcal{S}^{\pm}$. To see this, begin by letting $s \in \mathcal{S}^{\pm}$. Then $\sigma(s) = L_{\phi(1_G)^{-1}}\phi(s) = \phi(1_G)^{-1}\phi(s)$. Now, $e(1_G, s) \in E(\Gamma)$ and $\phi\big(e(1_G, s)\big) = e\big(\phi(1_G), \phi(1_G)^{-1}\phi(s)\big)$. Since ϕ is color preserving, $\phi(1_G)^{-1}\phi(s) = s'$, for some $s' \in \mathcal{S}^{\pm}$, or equivalently, $\phi(s) =$

$\phi(1_G)s'$. So, $\sigma(s) = \phi(1_G)^{-1}\phi(s) = \phi(1_G)^{-1}\phi(1_G)s' = s' \in \mathcal{S}^{\pm}$. Since s was chosen arbitrarily, we have shown that $\sigma(\mathcal{S}^{\pm}) = \mathcal{S}^{\pm}$.

We must now show that $\sigma \in \mathrm{Aut}(G)$. Since $\sigma \in \mathrm{Aut}_{\text{c.p.}}(\Gamma)$, we know that σ is a bijection on $V(\Gamma) = G$. We also need to show that σ preserves the group operation. To this end, let $g \in G$. Then $e(g,s) \in E(\Gamma)$ and $c\big(e(g,s)\big) = c\big(e(1_G,s)\big)$, so we must have $c\Big(\phi\big(e(g,s)\big)\Big) = c\Big(\phi\big(e(1_G,s)\big)\Big)$ since ϕ is color-preserving. From this we can deduce that $\phi(g)^{-1}\phi(gs) = \phi(1_G)^{-1}\phi(s)$. But $\phi(1_G)^{-1}\phi(s) = \sigma(s) = s'$ for some $s' \in \mathcal{S}^{\pm}$ since $\sigma(\mathcal{S}^{\pm}) = \mathcal{S}^{\pm}$, and so $\phi(g)^{-1}\phi(gs) = s'$, implying that $\phi(gs) = \phi(g)s'$. Then we can observe the following:

$$\sigma(gs) = L_{\phi(1_G)^{-1}}\phi(gs) = \phi(1_G)^{-1}\phi(gs) = \phi(1_G)^{-1}\phi(g)s' = \sigma(g)s' = \sigma(g)\sigma(s).$$

Thus, $\sigma(gs) = \sigma(g)\sigma(s)$ for any $s \in \mathcal{S}^{\pm}$ and for any $g \in G$. We must now show that $\sigma(gh) = \sigma(g)\sigma(h)$ for any $g,h \in G$. The element $h \in G$ can be written as a word in $\mathcal{S}^{\pm}$, since $\mathcal{S}$ generates G. Suppose $h = s_1 \cdots s_n$, where each $s_i \in \mathcal{S}^{\pm}$. Let's induct on n.

Suppose $h = s_1$. Then $\sigma(gh) = \sigma(gs_1) = \sigma(g)\sigma(s_1)$, by the preceding special case. So now assume that the result is true for some k and let's show it is true for $k+1$. Suppose $h = s_1 \cdots s_k s_{k+1}$. Then observe the following:

$$\begin{aligned}
\sigma(g)\sigma(h) &= \sigma(g)\sigma(s_1 \cdots s_k s_{k+1}) \\
&= \sigma(g)\sigma(g' s_{k+1}) \quad \text{(where } g' = s_1 \cdots s_k\text{)} \\
&= \sigma(g)\sigma(g')\sigma(s_{k+1}) \quad \text{(by the initial step)} \\
&= \sigma(g)\sigma(s_1 \cdots s_k)\sigma(s_{k+1})
\end{aligned}$$

$$= \sigma(gs_1 \cdots s_k)\sigma(s_{k+1}) \quad \text{(by inductive hypothesis)}$$

$$= \sigma(gs_1 \cdots s_k s_{k+1}) \quad \text{(by the initial step)}$$

$$= \sigma(gh).$$

Thus, the result is proved by induction, that is, for all $g, h \in G$, $\sigma(gh) = \sigma(g)\sigma(h)$. Therefore, $\sigma \in \mathrm{Aut}(G)$ and $\sigma(\mathcal{S}^{\pm}) = \mathcal{S}^{\pm}$. So, $\sigma \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$.

We now can see that $\phi \in G \cdot \mathrm{Aut}(G, \mathcal{S}^{\pm})$, since $\phi = L_{1_G}\gamma = L_{\phi(1_G)\phi(1_G)^{-1}}\phi = L_{\phi(1_G)}L_{\phi(1_G)^{-1}}\phi = L_{\phi(1_G)}\sigma$, where $L_{\phi(1_G)} \in G$ and $\sigma \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$.

Therefore, $\mathrm{Aut}_{\mathrm{c.p.}}(\Gamma) \subseteq G \cdot \mathrm{Aut}(G, \mathcal{S}^{\pm})$, and so $\mathrm{Aut}_{\mathrm{c.p.}}(\Gamma) = G \cdot \mathrm{Aut}(G, \mathcal{S}^{\pm})$.

Step 5: Summary

In conclusion, we've shown that G and $\mathrm{Aut}(G, \mathcal{S}^{\pm})$ are subgroups of $\mathrm{Aut}_{\mathrm{c.p.}}(\Gamma)$, with G a normal subgroup, $\mathrm{Aut}_{\mathrm{c.p.}}(\Gamma) = G \cdot \mathrm{Aut}(G, \mathcal{S}^{\pm})$, and $G \cap \mathrm{Aut}(G, \mathcal{S}^{\pm}) = \{\mathrm{id}_\Gamma\}$. Therefore, $\mathrm{Aut}_{\mathrm{c.p.}}(\Gamma) \cong G \rtimes \mathrm{Aut}(G, \mathcal{S}^{\pm})$.

□

Chapter 4

Reflections

4.1 Definitions, Basic Properties, and Examples

Throughout this chapter, we will assume that the graphs we work with are all connected. In the preceding chapter, we defined general graph morphisms. Here, we will focus on one special type of graph automorphism, called *reflection*. At the root of this type of automorphism is the concept of *inverted edges*. Given a graph automorphism $r : \Gamma \to \Gamma$ and an edge $e \in E(\Gamma)$, we say that the edge e is **inverted** if $r(e) = \overline{e}$. A **reflection** on a graph Γ is a graph automorphism $r : \Gamma \to \Gamma$ satisfying the following properties:

(i) $r^2 = \mathrm{id}_\Gamma$

(ii) the set of inverted edges $\Gamma_r = \{e \in E(\Gamma) \mid r(e) = \overline{e}\}$ separates Γ

When r is a reflection on Γ, the set Γ_r will be called the **wall of r** and the edges in Γ_r will be referred to as ***r*-reflectors**. Note that if $e \in \Gamma_r$, then $\overline{e} \in \Gamma_r$

also, since r must respect the reversing function. Therefore, $r(\overline{e}) = \overline{r(e)} = \overline{\overline{e}}$.

Example 4.1. Let $\Gamma = \Gamma(\mathbb{Z}, \mathcal{S})$, where $\mathbb{Z} = \langle a \rangle$ and $\mathcal{S} = \{a\}$. The graph Γ is depicted in Figure 4.1. Consider the function $\phi : \mathbb{Z} \to \mathbb{Z}$ defined by stating

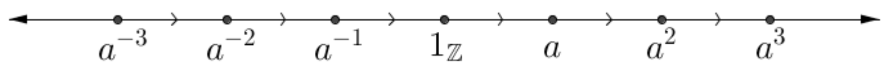

Figure 4.1: Cayley graph $\Gamma(\mathbb{Z}, \{a\})$

that $\phi(a) = a^{-1}$. Clearly $\phi \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$, and by Theorem 3.8 we know that $L_a\phi$ is a color-preserving graph automorphism on Γ. Let $e \in E(\Gamma)$. Then e must have the form $e = e(g, a)$, where $g \in \mathbb{Z}$. Since $\mathbb{Z}$ is generated by a, there exists an integer n such that $g = a^n$. So $e = e(a^n, a)$. Suppose e is inverted by $L_a\phi$. Then $L_a\phi(a^n) = a^n a$ and $L_a\phi(a^n a) = a^n$, and from this we can observe the following:

$$a^{n+1} = a^n a = L_a\phi(a^n) = a\left(\phi(a)\right)^n = a\left(a^{-1}\right)^n = a^{1-n}$$

$$a^n = L_a\phi\left(a^n a\right) = L_a\phi\left(a^{n+1}\right) = a\left(\phi(a)\right)^{n+1} = a\left(a^{-1}\right)^{n+1} = a^{1-n-1} = a^{-n}$$

Therefore, $a^{n+1} = a^{1-n}$ and $a^n = a^{-n}$, both of which imply that $a^{2n} = 1_{\mathbb{Z}}$. But this happens only when $n = 0$. So in order for $e = e(a^n, a) = a(g, a)$ to be inverted by $L_a\phi$, we must have $n = 0$ and $e = e(1_{\mathbb{Z}}, a)$. Then $e = e(1_{\mathbb{Z}}, a)$, together with $\overline{e} = e(a, a^{-1})$, are the only edges in Γ inverted by $L_a\phi$. The only edges in Γ between the vertices $1_{\mathbb{Z}}$ and a are e and $\overline{e}$, and so the set of inverted edges $\Gamma_{L_a\phi} = \{e, \overline{e}\} \subseteq E(\Gamma)$ separates the graph. Therefore, $L_a\phi$ is a reflection on Γ.

◇

Lemma 4.2. *Suppose r is a reflection on the graph Γ and γ is any nonempty path in Γ. If γ does not contain any r-reflectors, then $\iota(\gamma)$ and $\tau(\gamma)$ are in the same connected component of $\Gamma - \Gamma_r$.*

Proof. Assume γ is any nonempty path in Γ that does not contain any r-reflectors. By inducting on the length of the path γ, we will show that $\iota(\gamma)$ and $\tau(\gamma)$ are in the same connected component of $\Gamma - \Gamma_r$.

Since γ is a nonempty path, $\iota(\gamma)$ and $\tau(\gamma)$ are distinct vertices, and so γ must have length at least one. If γ has length 1, then γ consists of a single edge from $\iota(\gamma)$ to $\tau(\gamma)$. This single edge is not an r-reflector, since we are assuming that γ does not contain any r-reflectors. Therefore, this edge must be contained in a single connected component of $\Gamma - \Gamma_r$, implying that $\iota(\gamma)$ and $\tau(\gamma)$ are in the same connected component of $\Gamma - \Gamma_r$.

Now assume that if γ' is any path in Γ which does not contain any r-reflectors and if γ' has length k, then $\iota(\gamma')$ and $\tau(\gamma')$ are in the same connected component of $\Gamma - \Gamma_r$. Suppose γ is a path of length $k+1$ which does not contain any r-reflectors. Then there exist edges $e_1, \dots, e_{k+1}$ in $E(\Gamma)$ such that $\gamma = (e_1, \dots, e_{k+1})$. Consider the path $\gamma' = (e_1, \dots, e_k)$. Since γ does not contain any r-reflectors and since γ' is a subpath of γ, γ' also does not contain any r-reflectors. Then by the inductive hypothesis, $\iota(\gamma')$ and $\tau(\gamma')$ must be in the same connected component of $\Gamma - \Gamma_r$ since γ' has length k. But $\iota(e_{k+1})$ and $\tau(e_{k+1})$ must also be in the same connected component of $\Gamma - \Gamma_r$, by the initial step of the induction. Observe that $\iota(\gamma') = \iota(\gamma)$, $\tau(\gamma') = \tau(e_k) = \iota(e_{k+1})$, and $\tau(e_{k+1}) = \tau(\gamma)$. Since being in the same connected component of a graph is an equivalence relation, we can conclude by transitivity that $\iota(\gamma)$ and $\tau(\gamma)$ are in the same connected component of

$\Gamma - \Gamma_r$, completing the inductive proof.

□

Lemma 4.3. *Let r be a reflection on the graph Γ and let u and v be vertices in Γ. Suppose γ is a path in Γ from u to v. Then u and v are in the same connected component of $\Gamma - \Gamma_r$ if and only if γ contains an even number of r-reflectors.*

Proof. ($\Longleftarrow$) Assume that γ contains an even number of r-reflectors. If u and v are the same vertex, then they obviously are in the same connected component of $\Gamma - \Gamma_r$. So assume that $u \neq v$. If γ does not contain any r-reflectors, then Lemma 4.2 tells us that u and v are in the same connected component of $\Gamma - \Gamma_r$.

Suppose now that γ contains $2n$ r-reflectors, where $n \in \mathbb{N}$, and let e_1 and e_2 be two of these r-reflectors. Then there exist subpaths γ_1, γ_2, and γ_3 of γ such that $\gamma = \gamma_1 e_1 \gamma_2 e_2 \gamma_3$. Since e_1 and e_2 are r-reflectors, $\iota(r(\gamma_2)) = r(\iota(\gamma_2)) = r(\tau(e_1)) = \iota(e_1) = \tau(\gamma_1)$ and $\tau(r(\gamma_2)) = r(\tau(\gamma_2)) = r(\iota(e_2)) = \tau(e_2) = \iota(\gamma_3)$. Therefore, $\gamma_1 r(\gamma_2)\gamma_3$ is a well-defined path in Γ from u to v which does not contain edges e_1 or e_2. Figure 4.2 illustrates the scenario just described.

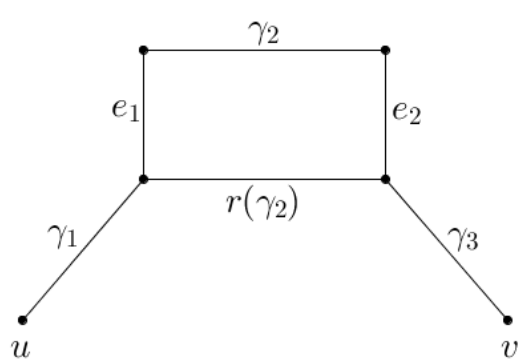

Figure 4.2: Eliminating pairs of r-reflectors to shorten a path

The path $\gamma_1 r(\gamma_2)\gamma_3$ contains $2n-2$ r-reflectors. By eliminating r-reflectors pairwise in this fashion, we will eventually have a path γ' in Γ from u to v which does not contain any r-reflectors. Lemma 4.2 then tells us that u and v are in the same connected component of $\Gamma - \Gamma_r$, as desired.

($\Longrightarrow$) We will prove this implication by proving the contrapositive of the implication. To this end, assume γ contains an odd number of r-reflectors. We will show that u and v must be in different connected components of $\Gamma - \Gamma_r$. Since $\Gamma - \Gamma_r$ is disconnected, there exists a path γ' in Γ containing exactly one r-reflector with initial vertex v and terminal vertex, say w, which lies in a different connected component of $\Gamma - \Gamma_r$ than v. Consider the path $\gamma\gamma'$ which has initial vertex u and terminal vertex w. This path has exactly one more r-reflector than path γ, meaning $\gamma\gamma'$ contains an even number of r-reflectors. Our previous work then tells us that u and w must be in the same connected component of $\Gamma - \Gamma_r$, implying that u and v are in different connected components of $\Gamma - \Gamma_r$. This completes the proof.

□

Corollary 4.4. *If r is a reflection on Γ, then any cycle in Γ must contain an even number of r-reflectors.*

Proof. Suppose that γ is a cycle in Γ based at a vertex v_0. Clearly $\iota(\gamma)$ and $\tau(\gamma)$ are in the same connected component of $\Gamma - \Gamma_r$, since $\iota(\gamma) = \tau(\gamma)$. Then Lemma 4.3 tells us that the cycle γ must contain an even number of r-reflectors.

□

Corollary 4.5. *The endpoints of any r-reflector are in different connected components of $\Gamma - \Gamma_r$.*

Proof. If e is an r-reflector, then e itself is a path between $\iota(e)$ and $\tau(e)$. This path obviously contains one r-reflector, and Lemma 4.3 tells us that $\iota(e)$ and $\tau(e)$ are in different connected components of $\Gamma - \Gamma_r$.

□

Corollary 4.6. *Let r be a reflection on Γ. If γ is a triangle in Γ, then none of the edges in γ are r-reflectors.*

Proof. Assume $\gamma = (e_1, e_2, e_3)$ is a triangle in Γ. Suppose γ contains an r-reflector. Then it must contain two r-reflectors since, according to Corollary 4.4, any cycle in Γ must contain an even number of r-reflectors. Without loss of generality, assume e_1 and e_2 are the r-reflectors in γ. Since r is a reflection, we must have $r(\tau(e_1)) = \iota(e_1)$ and $r(\iota(e_2)) = \tau(e_2) = \iota(e_3)$. However, $\tau(e_1) = \iota(e_2)$, and so we must have $\iota(e_1) = r(\tau(e_1)) = r(\iota(e_2)) = \iota(e_3)$. But $\iota(e_1)$ and $\iota(e_3)$ cannot be equal, since γ is a cycle. Thus, we have a contradiction! Therefore, the triangle γ cannot contain any r-reflectors.

□

The preceding proof hinges on the idea that two r-reflectors cannot share an endpoint. This is true in a general sense–not just for edges in a triangle. This small fact will sometimes be useful when proving later results involving reflections. Lemma 4.3 and Corollaries 4.4, 4.5, and 4.6 lay out some of the other basic principles that we will use when working with reflections. For instance, if r is a reflection on Γ, we now know that *any* path between the endpoints of an r-reflector must contain at least one r-reflector. This property gives us a basic way of showing that a graph automorphism is not a reflection. If ϕ is a graph automorphism on Γ which inverts an edge $e \in E(\Gamma)$

and if there exists a path γ in Γ from $\iota(e)$ to $\tau(e)$ such that ϕ does not invert any edges of γ, then ϕ cannot be a reflection.

The word *reflection* brings to mind thoughts of symmetry. The next lemma will show that a reflection on a graph Γ does possess strong symmetrical behavior.

Lemma 4.7. *Suppose r is a reflection on a graph Γ. The subgraph $\Gamma - \Gamma_r$ has exactly two connected components. Moreover, these two connected components are interchanged by r.*

Proof. By definition of reflection, $\Gamma - \Gamma_r$ is disconnected. Therefore, $\Gamma - \Gamma_r$ contains at least two connected components. Choose two vertices u and v in Γ which lie in distinct connected components of $\Gamma - \Gamma_r$. Since we are assuming that Γ is connected, there exists a path γ in Γ between u and v. By Lemma 4.3, γ must contain an odd number of r-reflectors. Let w be any other point in Γ distinct from u and v and let γ' be any path in Γ from v to w. Consider the path $\gamma\gamma'$ from u to w in Γ. If γ' contains an even number of r-reflectors, then by Lemma 4.3 we can conclude that v and w are in the same connected component of $\Gamma - \Gamma_r$. If γ' contains an odd number of r-reflectors, then the path $\gamma\gamma'$ contains an even number of r-reflectors and Lemma 4.3 tells us that u and w are in the same connected component of $\Gamma - \Gamma_r$. Therefore, there must be only two connected components of $\Gamma - \Gamma_r$, say Γ_1 and Γ_2.

Suppose u is a vertex in Γ. Without loss of generality we may assume that $u \in \Gamma_1$. Let $e \in E(\Gamma)$ be an arbitrary r-reflector and let v be the endpoint of e which lies in Γ_1. Then there exists a path γ in Γ between u and v which lies completely in Γ_1. Consequently, γ does not contain any

r-reflectors. Since r is a graph automorphism which sends each edge of Γ_r to its reverse, r sends non-r-reflectors to other non-r-reflectors. Moreover, as a graph automorphism, r preserves paths. Since γ is a path in Γ from u to v, $r(\gamma)$ is a path in Γ from $r(u)$ to $r(v)$. Since γ does not contain any r-reflectors, $r(\gamma)$ does not contain any r-reflectors. Therefore, $r(u)$ and $r(v)$ must be in the same connected component of $\Gamma - \Gamma_r$. But since e is an r-reflector and v is an endpoint of e which lies in Γ_1, $r(v)$ must lie in Γ_2. Thus, $r(u)$ must lie in Γ_2 also. Since r is a graph automorphism, we now can conclude that $r(\Gamma_1) = \Gamma_2$ and $r(\Gamma_2) = \Gamma_1$

Therefore, the reflection r interchanges the connected components of $\Gamma - \Gamma_r$.

□

Corollary 4.8. *If r is a reflection on Γ, then r is* ***vertex free on Γ****, meaning that for any $v \in V(\Gamma)$ we have $r(v) \neq v$.*

In general, it can be difficult to show that a graph automorphism ϕ on Γ is a reflection. First, one must identify the set Γ_ϕ of edges in Γ which are inverted by ϕ. Second, one must show that the set Γ_ϕ separates the graph. The next lemma gives us a way of showing that ϕ is a reflection by first identifying a set $\mathcal{E}$ of edges which separates Γ and then proving that all edges in $\mathcal{E}$ are inverted by ϕ.

Lemma 4.9. *Let ϕ be an automorphism on a graph Γ such that $\phi^2 = id_\Gamma$. Suppose $\mathcal{E} \subseteq E(\Gamma)$ is a collection of edges which separates Γ. If ϕ inverts every edge in $\mathcal{E}$, then ϕ is a reflection on Γ and $\Gamma_\phi = \mathcal{E}$.*

Proof. Assume ϕ inverts every edge in $\mathcal{E}$. Then $\mathcal{E}$ must be contained in Γ_ϕ, the set of all edges in Γ which are inverted by ϕ. By assumption, $\Gamma - \mathcal{E}$ is disconnected. But $\mathcal{E} \subseteq \Gamma_\phi$ implies that $\Gamma - \Gamma_\phi = (\Gamma - \mathcal{E}) - (\Gamma_\phi - \mathcal{E})$. In other words, if $\mathcal{E} = \Gamma_\phi$, then $\Gamma - \Gamma_\phi = \Gamma - \mathcal{E}$, and if $\mathcal{E} \neq \Gamma_\phi$, then $\Gamma - \Gamma_\phi$ can be obtained from $\Gamma - \mathcal{E}$ by removing additional edges. In either event, we can see that $\Gamma - \Gamma_\phi$ must be disconnected, since $\Gamma - \mathcal{E}$ is disconnected. Therefore, ϕ is an automorphism on Γ with $\phi^2 = \mathrm{id}_\Gamma$ such that Γ_ϕ, the set of all edges in Γ which are inverted by ϕ, separates Γ. Thus, ϕ is a reflection on Γ.

Lemma 4.7 tells us that $\Gamma - \Gamma_\phi$ has exactly two connected components. Call these components Γ_1 and Γ_2. We now will show that $\Gamma_\phi = \mathcal{E}$. We already know that $\mathcal{E} \subseteq \Gamma_\phi$. Suppose $e \in \Gamma_\phi$ and consider the set of edges $\Gamma_0 \coloneqq \Gamma_\phi - \{e, \overline{e}\}$. The graph $\Gamma - \Gamma_0$ consists of disjoint subgraph Γ_1 and Γ_2, along with the edge pair $\{e, \overline{e}\}$. We know that $\Gamma - \Gamma_\phi$ is disconnected, but $\Gamma - \Gamma_0$ must be connected, since the edge e serves as a bridge between the subgraphs Γ_1 and Γ_2. We then can observe that if $\widetilde{\Gamma}$ is any proper subset of Γ_ϕ, then $\Gamma - \widetilde{\Gamma}$ will be connected. Consequently, since $\mathcal{E}$ separates Γ, $\mathcal{E}$ cannot be a proper subset of Γ_ϕ. In conclusion, $\mathcal{E} = \Gamma_\phi$.

□

The following lemma prescribes a way of creating new reflections from already known reflections.

Lemma 4.10. *If r is a reflection on a graph Γ and ϕ is any automorphism on Γ, then $\phi r \phi^{-1}$ is a reflection on Γ. Moreover, $\Gamma_{\phi r \phi^{-1}} = \phi(\Gamma_r)$.*

Proof. First we will show that if $\mathcal{E} \subseteq E(\Gamma)$ is a collection of edges which separates Γ, then $\phi(\mathcal{E})$ also separates Γ. To see this, let $e \in \phi(\mathcal{E})$. Then there exists $e' \in \mathcal{E}$ such that $\phi(e') = e$. Let γ be any path in Γ from $\iota(e)$ and $\tau(e)$. Then $\phi^{-1}(\gamma)$ is a path in Γ from $\phi^{-1}(\iota(e))$ to $\phi^{-1}(\tau(e))$. Because

ϕ^{-1} is an automorphism on Γ, it must respect the initial-vertex and reversing functions inherent to Γ, which allows us to observe the following:

$$\phi^{-1}\big(\iota(e)\big) = \iota\left(\phi^{-1}(e)\right) = \iota\left(\phi^{-1}\left(\phi(e')\right)\right) = \iota(e')$$
$$\phi^{-1}\big(\iota(\overline{e})\big) = \iota\big(\phi^{-1}(\overline{e})\big) = \iota\big(\overline{\phi^{-1}(e)}\big) = \tau\left(\phi^{-1}(e)\right) = \tau\left(\phi^{-1}\left(\phi(e')\right)\right) = \tau(e')$$

Therefore, $\phi^{-1}(\gamma)$ is a path in Γ from $\iota(e')$ to $\tau(e')$. Since $e' \in \mathcal{E}$ and $\mathcal{E}$ separates Γ, $\phi^{-1}(\gamma)$ must contain at least one edge from $\mathcal{E}$, say $\widetilde{e}$. But then we can see that $\phi(\widetilde{e}) \in \phi(\mathcal{E})$ and $\phi(\widetilde{e})$ is an edge in $\phi\left(\phi^{-1}(\gamma)\right) = \gamma$. Since the path γ from $\iota(e)$ to $\tau(e)$ was arbitrary, we can now conclude that any path in Γ from $\iota(e)$ to $\tau(e)$ must contain an edge from $\phi(\mathcal{E})$. Since $e \in \phi(\mathcal{E})$ was arbitrary, we can now conclude that $\phi(\mathcal{E})$ separates Γ, as desired.

Now, suppose $\mathcal{E} := \phi(\Gamma_r)$. Since Γ_r separates Γ and ϕ is an automorphism on Γ, our preceding work allows us to conclude that $\mathcal{E}$ separates Γ. Since $r^2 = \mathrm{id}_\Gamma$, we can easily observe that $\left(\phi r \phi^{-1}\right)^2 = \mathrm{id}_\Gamma$. If we can show that $\phi r \phi^{-1}$ inverts every edge in $\mathcal{E}$, then Lemma 4.9 will tell us that $\phi r \phi^{-1}$ is a reflection on Γ and, moreover, that $\Gamma_{\phi r \phi^{-1}} = \mathcal{E} = \phi(\Gamma_r)$. Let $e \in \mathcal{E}$. Then there exists $e' \in \Gamma_r$ such that $e = \phi(e')$. Since e' is an r-reflector, we must have $r(\iota(e')) = \tau(e')$ and $r(\tau(e')) = \iota(e')$. Then observe the following:

$$\phi r \phi^{-1}(\iota(e)) = \phi r\left(\iota\left(\phi^{-1}(e)\right)\right) = \phi r(\iota(e')) = \phi(\tau(e')) = \tau(\phi(e')) = \tau(e)$$
$$\phi r \phi^{-1}(\tau(e)) = \phi r\left(\tau\left(\phi^{-1}(e)\right)\right) = \phi r(\tau(e')) = \phi(\iota(e')) = \iota(\phi(e')) = \iota(e)$$

Therefore, $\phi r \phi^{-1}$ inverts the edge e. Since $e \in \mathcal{E}$ was arbitrary, we now know that $\phi r \phi^{-1}$ inverts every edge in $\mathcal{E}$, thereby completing the proof.

□

Example 4.11. Consider again the graph $\Gamma = \Gamma(G, \mathcal{S})$, where $G = \mathbb{Z} = \langle a \rangle$ and $\mathcal{S} = \{a\}$. We showed in Example 4.1 that $L_a\phi$ is a reflection on Γ, where $\phi : \mathbb{Z} \to \mathbb{Z}$ is defined by stating that $\phi(a) = a^{-1}$. Additionally, we showed that the wall of $L_a\phi$ is $\Gamma_{L_a\phi} = \{e, \overline{e}\}$, where $e = e(1_{\mathbb{Z}}, a)$.

Lemma 3.6 of Chapter 3 tells us that L_g is a graph automorphism on Γ for every $g \in \mathbb{Z}$. Then Lemma 4.10 that we just proved tells us that $L_g L_a \phi (L_g)^{-1} = L_g L_a \phi L_{g^{-1}}$ also is a reflection on Γ and $\Gamma_{L_g L_a \phi L_{g^{-1}}} = L_g(\Gamma_{L_a\phi})$. We can easily see that $L_g(e) = L_g(e(1_{\mathbb{Z}}, a)) = e(g, a)$. Therefore, $\Gamma_{L_g L_a \phi L_{g^{-1}}} = \{e(g,a), \overline{e(g,a)}\}$.

◇

4.2 Reflection Groups

If G is a subgroup of $\mathrm{Aut}(\Gamma)$, then the inclusion function from G to $\mathrm{Aut}(\Gamma)$ induces a group action of G on Γ. A subgroup G of $\mathrm{Aut}(\Gamma)$ is called a **reflection group on Γ** if G is generated by reflections on Γ and if the action of G on Γ is edge free. Let $\mathfrak{R}$ denote the set of all reflections contained in the reflection group G. From Lemma 4.10 we know that $\mathfrak{R}$ is closed under conjugation. We can also observe that if r and r' are distinct elements of $\mathfrak{R}$, then their walls Γ_r and $\Gamma_{r'}$ are disjoint. If $\Gamma_r \cap \Gamma_{r'} \neq \varnothing$, then there would be an edge e in Γ which is both an r-reflector and r'-reflector. Then rr' is in G and $rr'(e) = e$, implying that $rr' = \mathrm{id}_\Gamma$ since the action of G on Γ is edge free. But this implies that $r = r'$, contradicting our assumption that r and r' are distinct.

If G is a reflection group on Γ with reflection set $\mathfrak{R}$, then the subgraph

$\Gamma - \left(\bigcup_{r\in\mathfrak{R}} \Gamma_r\right)$ of Γ is disconnected since $\Gamma - \Gamma_r$ is disconnected for each $r \in \mathfrak{R}$. The connected components of $\Gamma - \left(\bigcup_{r\in\mathfrak{R}} \Gamma_r\right)$ are called the **chambers** of the reflection group G. If C is a chamber of G and if $r \in \mathfrak{R}$, we say that Γ_r is a **wall of C** if there exists an r-reflector e with $\iota(e)$ in C. Fix a vertex $v_0 \in V(\Gamma)$ which will serve as a base vertex. The unique chamber of G which contains v_0 is called the **fundamental chamber**, denoted C_{v_0}. If $r \in \mathfrak{R}$ and if Γ_r is a wall of C_{v_0}, then r will be referred to as a **fundamental reflection**. Suppose $\mathcal{S}$ is a set of fundamental reflections. Then the pair $(G, \mathcal{S})$ is called a **reflection system** on Γ.

Example 4.12. Let $H = \langle a, b \mid b^3 = 1_H, ab = ba\rangle \cong \mathbb{Z} \times \mathbb{Z}_3$ and let $\mathcal{T} = \{a, b\}$. Consider the Cayley graph $\Gamma = \Gamma(H, \mathcal{T})$ which is depicted in Figure 4.3.

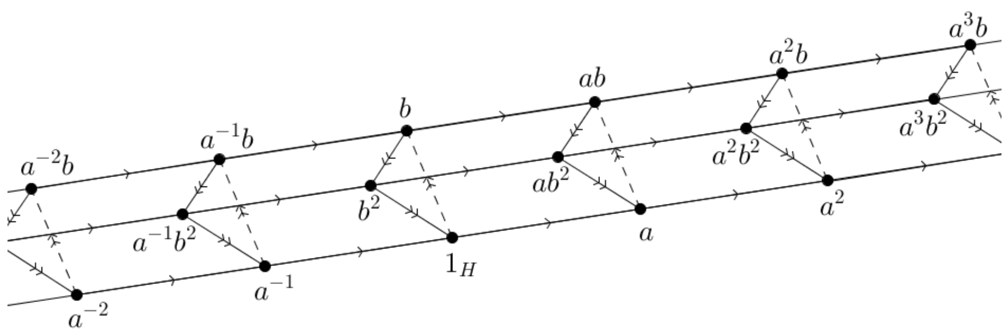

Figure 4.3: Cayley graph $\Gamma(\mathbb{Z} \times \mathbb{Z}_3, \{a, b\})$

Define a homomorphism $\phi : H \to H$ by stating that $\phi(a) = a^{-1}$ and $\phi(b) = b$. It is clear that ϕ respects the relations of H and thus is an automorphism of H. Moreover, $\phi(\mathcal{T}^{\pm}) = \mathcal{T}^{\pm}$. Therefore, $\phi \in \mathrm{Aut}(H, \mathcal{T}^{\pm})$. Then we know by Theorem 3.8 that $L_a\phi$ and $L_{a^{-1}}\phi$ are color-preserving graph automorphisms on Γ. One can check that $(L_a\phi)^2 = \mathrm{id}_\Gamma$ and $(L_{a^{-1}}\phi)^2 = \mathrm{id}_\Gamma$. The set of edges in Γ which are inverted by $L_a\phi$ is $\Gamma_{L_a\phi} = \{e_1 = e(1_H, a), e_2 = e(b, a), e_3 =$

$e(b^2, a), \overline{e_1}, \overline{e_2}, \overline{e_3}\}$, which separates Γ. Therefore, $L_a\phi$ is a reflection on Γ. We can also see that $L_{a^{-1}}\phi$ is a reflection on Γ with $\Gamma_{L_{a^{-1}}\phi} = \{e_4 = e(a^{-1}, a), e_5 = e(ba^{-1}, a), e_6 = e(b^2a^{-1}, a), \overline{e_4}, \overline{e_5}, \overline{e_6}\}$.

Let $G := \langle r_1, r_2 \rangle$, where $r_1 = L_a\phi$ and $r_2 = L_{a^{-1}}\phi$. Since r_1 and r_2 are both reflections on Γ, we will be able to conclude that G is a reflection group on Γ if we can show that the action of G on Γ is edge free. In order to determine this, we need to better understand the nature of the elements of G. Observe that $r_1r_2 = L_a\phi L_{a^{-1}}\phi = L_{a^2}$ and $r_2r_1 = L_{a^{-1}}\phi L_a\phi = L_{a^{-2}}$. Therefore, r_1 and r_2 have order 2 and do not commute, and so each element of G can be written as an alternating word in r_1 and r_2. One can check that the following observations are true for any nonnegative integer n:

$$(r_1r_2)^n = L_{a^{2n}}$$

$$(r_2r_1)^n = L_{a^{-2n}}$$

$$(r_1r_2)^n r_1 = L_{a^{2n+1}}\phi$$

$$(r_2r_1)^n r_2 = L_{a^{-(2n+1)}}\phi$$

Therefore, $G = \{L_{a^{2k}} \mid k \in \mathbb{Z}\} \cup \{L_{a^{2k+1}}\phi \mid k \in \mathbb{Z}\}$.

When $k \neq 0$, the elements of the form $L_{a^{2k}}$ act on Γ by translation. Moreover, these elements neither fix nor invert any edges in Γ. For any k, the elements of the form $L_{a^{2k+1}}\phi$ are reflections on Γ. For a fixed k, the wall of $L_{a^{2k+1}}\phi$ is as shown here:

$$\Gamma_{L_{a^{2k+1}}\phi} = \{e_1 = e(a^k, a), e_2 = e(ba^k, a), e_3 = e(b^2a^k, a), \overline{e_1}, \overline{e_2}, \overline{e_3}\}$$

We can see also that the elements of the form $L_{a^{2k+1}}\phi$ do not fix any edges in Γ. We can now conclude that id_Γ is the only element of G which fixes any edges in Γ. Therefore, G is a reflection group on Γ.

The set of all reflections in G is given by $\mathfrak{R} = \{L_{a^{2k+1}}\phi \mid k \in \mathbb{Z}\}$. Also, $\bigcup_{r\in\mathfrak{R}} \Gamma_r = \{e(h,a), e(h,a^{-1}) \mid h \in H\}$. The subgraph $\Gamma - \left(\bigcup_{r\in\mathfrak{R}} \Gamma_r\right)$ of Γ is shown here:

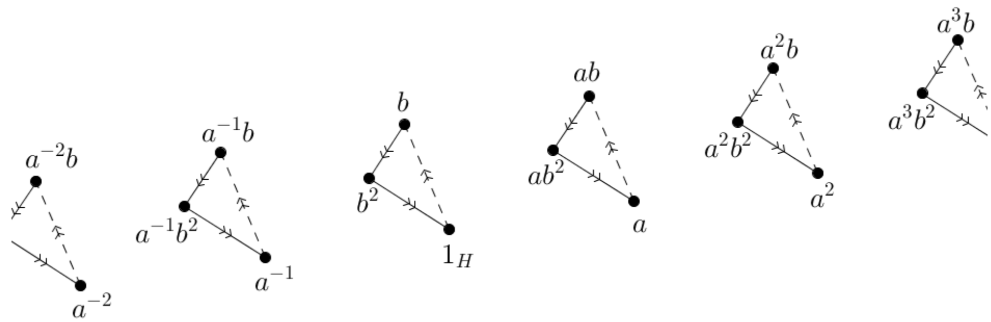

Figure 4.4: Chamber subgraph $\Gamma - \left(\bigcup_{r\in\mathfrak{R}} \Gamma_r\right)$ of $\Gamma(\mathbb{Z}\times\mathbb{Z}_3, \{a,b\})$

We can see clearly that the chambers of G correspond to the triangles shown above. Fix the vertex 1_H in Γ to be the base vertex. Then the fundamental chamber is $C_{1_H} = \{e_1 = e(1_H, b), e_2 = e(b,b), e_3 = e(b^2,b), \overline{e_1}, \overline{e_2}, \overline{e_3}\}$. The set of fundamental reflections is given by $\mathcal{S} = \{r_1, r_2\}$. Therefore, $(G,\mathcal{S})$ is a reflection system on Γ.

◇

4.3 Coxeter Groups

Recall that a **Coxeter group**, often denoted by the symbol W, is a group with a presentation of the form shown here:

$$W := \langle r_1, r_2, \ldots, r_n \mid (r_i r_j)^{m_{i,j}} = 1_W \rangle$$

where for $i, j \in I_n$, $m_{i,j} \in \mathbb{N} \cup \{\infty\}$; $m_{i,j} = m_{j,i}$; $m_{i,i} = 1_W$; and $m_{i,j} \geq 2$ for $i \neq j$. When $m_{i,j} = \infty$, we interpret this to mean that $(r_i r_j)^m \neq 1_W$ for any $m \in \mathbb{N}$. The condition that $m_{i,i} = 1_W$ implies that $(r_i)^2 = 1_W$ for all $i \in I_n$. The generating set $\mathcal{S} = \{r_1, \ldots, r_n\}$ often is called the **fundamental generating set for W**. The pair $(W, \mathcal{S})$ is referred to as a **Coxeter system**. Though it is not immediately clear from the group presentation, it can be shown that $m_{i,j}$ is, in fact, the order of $r_i r_j$ in W [8].

Coxeter groups have a rich history and a strong relationship with reflection groups. The theory is developed thoroughly in [4], [8], and [9]. We will mention a few results here which will be of use in later chapters or which will serve as important examples.

Lemma 4.13. *Suppose W is a Coxeter group with fundamental generating set $\mathcal{S}$. If $r \in \mathcal{S}$, then r is a reflection acting on $\Gamma = \Gamma(W, \mathcal{S})$ by left translation. Furthermore, the pair $(W, \mathcal{S})$ is a reflection system on Γ.*

Example 4.14. Suppose P_n is a regular polygon with n sides. Recall that for any positive integer n, the dihedral group D_n is defined in terms of symmetries on P_n. The $2n$ elements of D_n are as described here:

- the rotation symmetries on P_n, of which there are n
- the reflection symmetries on P_n, of which there are n

The group operation is taken to be composition. If t represents counterclockwise rotation by an angle of $\frac{2\pi}{n}$ on P_n and r represents one of the reflections on P_n, then we know that D_n can be given by the presentation shown here:

$$D_n = \langle t, r, \mid t^n = r^2 = (rt)^2 = 1_{D_n} \rangle$$

However, it is well known that D_n can be given another presentation with generators r_1 and r_2, where $r_1 = r$ and $r_2 = tr$, as shown here:

$$D_n = \langle r_1, r_2 \mid r_1^2 = r_2^2 = (r_1 r_2)^n = 1_{D_n} \rangle$$

This second presentation makes it clear that the dihedral group D_n is, in fact, a Coxeter group. If $\mathcal{S} = \{r_1, r_2\}$, then Lemma 4.13 tells us that $(D_n, \mathcal{S})$ is a reflection system on $\Gamma(D_n, \mathcal{S})$. Interestingly, the Cayley graph $\Gamma(D_n, \mathcal{S})$ is a $2n$-sided polygon where the sides alternate colors between r_1 and r_2.

◇

Lemma 4.15. *Let $(W, \mathcal{S})$ be a reflection system on a graph Γ.*

(a) The action of W on Γ is simply transitive on the set of chambers.

(b) Every reflection in W is conjugate to a fundamental reflection.

(c) The set of fundamental reflections $\mathcal{S}$ generates the group W.

Theorem 4.16. *Let $(W, \mathcal{S})$ be a reflection system on a graph Γ. Then $(W, \mathcal{S})$ forms a Coxeter system, where for any $r_i, r_j \in \mathcal{S}$, $m_{i,j}$ is taken to be the order of the element $r_i r_j$ in W.*

Chapter 5

Totally Reflected Group Systems

5.1 Definitions, Basic Properties, and Examples

The theory that we developed in Chapter 4 concerning reflections applied to generic graphs. However, the examples that we used all involved Cayley graphs. It will soon become clear that the attention we paid to groups and their Cayley graphs in earlier chapters was intended to set the stage for this chapter.

One group may have many different generating sets. If $\mathcal{S}$ and $\mathcal{T}$ are two different generating sets for G, then the Cayley graphs $\Gamma(G, \mathcal{S})$ and $\Gamma(G, \mathcal{T})$ can vary remarkably. We gave an explicit example of this phenomenon in Chapter 2 involving the dihedral group D_4. Figures 2.1 and 2.2 depict Cayley graphs for D_4 with respect to two different generating sets. If we were to define a property that a given graph might (or might not) satisfy, we would not be surprised to find that there exist generating sets $\mathcal{S}$ and $\mathcal{T}$ for G such that $\Gamma(G,\mathcal{S})$ satisfies the property but $\Gamma(G,\mathcal{T})$ does not. We will now define such a property.

If G is a (nontrivial) group with (nice) generating set $\mathcal{S}$, then we will refer to the pair $(G,\mathcal{S})$ as a **group system**. We will say that the group system $(G,\mathcal{S})$ is **totally reflected** (or is a **totally reflected system**) if for each edge e in $\Gamma = \Gamma(G,\mathcal{S})$ there exists a color-preserving graph reflection on Γ which inverts e. Additionally, we will say that the group G is **totally reflected with respect to $\mathcal{S}$** if the system $(G,\mathcal{S})$ is totally reflected.

In previous chapters we explored the notions of graph reflections and color-preserving graph automorphisms. These concepts come together in the definition of *totally reflected systems*. If G is a group with generating set $\mathcal{S}$, then Theorem 3.8 tells us that any color-preserving automorphism on $\Gamma = \Gamma(G,\mathcal{S})$ has the form $L_k\alpha$, where $k \in G$ and $\alpha \in \mathrm{Aut}(G,\mathcal{S}^{\pm})$. If we assume also that $L_k\alpha$ is a reflection on Γ, then we must have that $(L_k\alpha)^2 = \mathrm{id}_\Gamma$. Therefore, for any $x \in G$, the following must be true:

$$x = (L_k\alpha)^2(x) = L_k\alpha\left(L_k\alpha(x)\right) = k\alpha(k)\alpha^2(x)$$

In particular, $1_G = k\alpha(k)\alpha^2(1_G) = k\alpha(k)$, implying that $\alpha(k) = k^{-1}$. Then we can conclude that $x = k\alpha(k)\alpha^2(x) = kk^{-1}\alpha^2(x) = \alpha^2(x)$ for all $x \in G$. Consequently, $\alpha^2 = \mathrm{id}_G$.

If $L_k\alpha$ inverts an edge $e(g,s)$ in Γ, where $g \in G$ and $s \in \mathcal{S}^{\pm}$, then we can make additional observations concerning $L_k\alpha$, as shown here:

$$gs = L_k\alpha(g) = k\alpha(g)$$

$$g = L_k\alpha(gs) = k\alpha(g)\alpha(s) = gs\alpha(s)$$

If we multiply on the left by g^{-1} in the right-hand equation, we can see that $1_G = s\alpha(s)$ and so $\alpha(s) = s^{-1}$. Let us summarize these observations in the next lemma.

Lemma 5.1. *Let $L_k\alpha$ be a color-preserving automorphism on a Cayley graph $\Gamma = \Gamma(G,\mathcal{S})$, where $k \in G$ and $\alpha \in \mathrm{Aut}(G,\mathcal{S}^{\pm})$. If $L_k\alpha$ is a reflection on Γ, then α must satisfy the extra conditions that $\alpha(k) = k^{-1}$ and $\alpha^2 = id_G$. Moreover, if $L_k\alpha$ is a reflection on Γ which inverts an edge of the form $e(g,s)$, then we must also have that $k\alpha(g) = gs$ and $\alpha(s) = s^{-1}$.*

Corollary 5.2. *For any $s \in \mathcal{S}^{\pm}$, a color-preserving reflection on $\Gamma = \Gamma(G,\mathcal{S})$ which inverts the edge $e(1_G,s)$ must be of the form $L_s\alpha$, where $\alpha \in \mathrm{Aut}(G,\mathcal{S}^{\pm})$ satisfies $\alpha(s) = s^{-1}$ and $\alpha^2 = id_G$.*

The characterizations of color-preserving reflections on a Cayley graph which Lemma 5.1 and Corollary 5.2 provide will be very useful in our proofs involving totally reflected systems. If a group G has other nice properties, then it may be possible to give a more detailed characterization of the color-preserving reflections on a Cayley graph for G.

Lemma 5.3. *Let G be a group with generating set $\mathcal{S}$. Suppose $L_k\alpha$ is a color-preserving reflection on $\Gamma = \Gamma(G,\mathcal{S})$ which inverts the edge $e(g,s)$, where $g \in G$ and $s \in \mathcal{S}^{\pm}$. If $t \in \mathcal{S}^{\pm} - \{s,s^{-1}\}$ and if t commutes with s, then we must have $\alpha(t) = t$.*

Proof. Let t be an arbitrary element of $\mathcal{S}^{\pm} - \{s, s^{-1}\}$ such that $st = ts$. Then notice that $s^{-1}t = ts^{-1}$ as well. Let e_1, e_2, e_3, and e_4 be edges in Γ with $e_1 = e(g, s)$, $e_2 = e(gs, t)$, $e_3 = e(gst, s^{-1})$, and $e_4 = e(gsts^{-1}, t^{-1})$. Define a path γ in Γ by $\gamma = (e_1, e_2, e_3, e_4)$. We can easily see that γ is a well-defined path. Observe that the word in $\mathcal{S}^{\pm}$ corresponding to γ is $w(\gamma) = sts^{-1}t^{-1} = ss^{-1}tt^{-1} = 1_G$. Therefore, $\tau(e_4) = gsts^{-1}t^{-1} = g = \iota(e_1)$, and so γ is a closed path. Furthermore, since $s \neq t$ and $s \neq t^{-1}$, we must have that $\iota(e_1), \iota(e_2), \iota(e_3)$, and $\iota(e_4)$ are all distinct, implying that γ is a cycle in Γ. Since $L_k\alpha$ inverts the edge e_1, it must also invert the edge e_3. Using this fact along with the properties of $L_k\alpha$ established by Lemma 5.1, we can observe the following:

$$gt = gtss^{-1} = gsts^{-1} = L_k\alpha(gst) = k\alpha(g)\alpha(s)\alpha(t) = gss^{-1}\alpha(t) = g\alpha(t)$$

Then $gt = g\alpha(t)$ implies that $t = \alpha(t)$, which completes the proof.

□

The definition of *totally reflected system* leads to a natural question. For a group G, can there exist different generating sets $\mathcal{S}$ and $\mathcal{T}$ such that $(G, \mathcal{S})$ is totally reflected but $(G, \mathcal{T})$ is not? In general, the answer is "yes". However, we will show that there is at least one group which is not totally reflected with respect to *any* generating set. Likewise, we will show that there is exactly one nontrivial group which is totally reflected with respect to *every* generating set.

Example 5.4. $(\mathbb{Z}, \mathcal{S})$ is a totally reflected system, where $\mathbb{Z} = \langle a \rangle$ and $\mathcal{S} = \{a\}$. To see this, suppose $e(g, s)$ is an edge in $\Gamma = \Gamma(\mathbb{Z}, \{a\})$. Then $g = a^n$ for some integer n and $s \in \mathcal{S}^{\pm} = \{a, a^{-1}\}$. First, suppose $s = a$. Example 4.11 tells us that $r = L_{a^n} L_a \phi L_{a^{-n}}$ is a reflection on Γ which inverts the edge $e(g, s) = e(a^n, a)$, where $\phi : \mathbb{Z} \to \mathbb{Z}$ is the group automorphism defined by stating that $\phi(a) = a^{-1}$. We can easily see that $\phi(\mathcal{S}^{\pm}) = \mathcal{S}^{\pm}$. Therefore, $\phi \in \mathrm{Aut}(\mathbb{Z}, \mathcal{S}^{\pm})$. Let $x \in \mathbb{Z}$ be arbitrary. Observe the following:

$$L_{a^n} L_a \phi L_{a^{-n}}(x) = a^n a \phi(a^{-n} x) = a^{n+1} \phi(a^{-n}) \phi(x) = a^{n+1} a^n \phi(x) = L_{a^{2n+1}} \phi(x)$$

Then $r = L_{a^n} L_a \phi L_{a^{-n}} = L_{a^{2n+1}} \phi$, and Theorem 3.8 allows us to conclude that r is a color-preserving graph automorphism, in addition to being a reflection. Since g was arbitrary, we have shown that for any edge of the form $e(g, a)$ in Γ there exists a color-preserving graph reflection r on Γ which inverts $e(g, a)$.

Now suppose $s = a^{-1}$. Then $e(g, s) = e(a^n, a^{-1})$ is inverted by the color-preserving graph reflection $r = L_{a^{n-1}} L_a \phi L_{a^{1-n}} = L_{a^{2n-1}} \phi$, since our preceding work shows that r inverts the edge $\overline{e(a^n, a^{-1})} = e(a^{n-1}, a)$ in Γ.

Therefore, for any edge e in $\Gamma = \Gamma(\mathbb{Z}, \{a\})$, there exists a color-preserving graph reflection r on Γ which inverts e. Thus, $(\mathbb{Z}, \{a\})$ is a totally reflected system.

$\diamond$

The following proposition will simplify the matter of determining whether or not a pair $(G, \mathcal{S})$ is a totally reflected system.

Proposition 5.5. $(G, \mathcal{S})$ *is a totally reflected system if and only if for any* $s \in \mathcal{S}$ *there exists a color-preserving graph reflection* r *on* Γ *which inverts the edge* $e(1_G, s)$.

Proof. ($\Longrightarrow$) This direction is proved by the definition of totally reflected system.

($\Longleftarrow$) For any $t \in \mathcal{S}$, assume that there exists a color-preserving graph reflection, say r_t, on Γ which inverts the edge $e(1_G, t)$. For arbitrary $g \in G$ and $s \in \mathcal{S}$, consider the edge $e(g, s)$ in Γ. Let $r := L_g r_s L_{g^{-1}}$. Clearly r is a graph automorphism on Γ, since it is a composition of elements of $\mathrm{Aut}(\Gamma)$. Moreover, r is color-preserving, since $L_g, L_{g^{-1}} \in \mathrm{Aut}_{\text{c.f.}}(\Gamma) \subseteq \mathrm{Aut}_{\text{c.p.}}(\Gamma)$ and since $\mathrm{Aut}_{\text{c.p.}}(\Gamma) \leqslant \mathrm{Aut}(\Gamma)$ as shown by Lemma 3.7. Also, Lemma 4.10 tells us that r is a reflection on Γ. Since r_s inverts the edge $e(1_G, s)$, we must have $r_s(1_G) = s$ and $r_s(s) = 1_G$. Observe the following:

$$r(g) = L_g r_s L_{g^{-1}}(g) = g r_s(g^{-1} g) = g r_s(1_G) = gs$$
$$r(gs) = L_g r_s L_{g^{-1}}(gs) = g r_s(g^{-1} gs) = g r_s(s) = g 1_G = g$$

Therefore, r inverts the edge $e(g, s)$.

Now consider the edge $e(g, s^{-1})$ in Γ, where $g \in G$ and $s \in \mathcal{S}$ are arbitrary. Let $r := L_h r_s L_{h^{-1}}$, where $h = gs^{-1}$. From our preceding work, we know that r is a color-preserving graph reflection on Γ which inverts the edge $e(h, s)$. However, if $e(h, s)$ is an r-reflector, then so is $\overline{e(h, s)} = e(hs, s^{-1}) = e(gs^{-1}s, s^{-1}) = e(g, s^{-1})$. Therefore, r inverts the edge $e(g, s^{-1})$.

Any edge e in Γ is either of the form $e(g, s)$ or $e(g, s^{-1})$ for some $g \in G$ and $s \in \mathcal{S}$. We have shown, then, that for any edge e in Γ there exists a color-preserving graph reflection which inverts the edge e. Therefore, $(G, \mathcal{S})$ is a totally reflected system. $\square$

Example 5.6. Lemma 4.13 tells us that a Coxeter system $(W, \mathcal{S})$ is a reflection system on $\Gamma = \Gamma(W, \mathcal{S})$, where the action of W on Γ is taken to be left translation. Therefore, if $\mathcal{S} = \{r_1, \ldots, r_n\}$, then r_i is a reflection on Γ for any $i \in I_n$. One can easily check that r_i inverts the edge $e(1_W, r_i)$ in Γ. Also, r_i is color fixing and thus color preserving. Applying Proposition 5.5 then allows us to conclude that $(W, \mathcal{S})$ is a totally reflected system.

$\diamond$

Example 5.7. For any $n \in \mathbb{N}$, let $\mathbb{Z}_{2n} = \langle a \mid a^{2n} = 1_{\mathbb{Z}_{2n}} \rangle$ be the finite cyclic group of order $2n$. Suppose $\mathcal{S} = \{a\}$. Consider the color-preserving automorphism on $\Gamma = \Gamma(\mathbb{Z}_{2n}, \mathcal{S})$ given by $L_a\phi$, where $\phi \in \mathrm{Aut}(\mathbb{Z}_{2n}, \mathcal{S}^{\pm})$ is defined by stating that $\phi(a) = a^{-1}$. One can easily check that $L_a\phi$ is a reflection on Γ with wall $\Gamma_{L_a\phi} = \{e_1 = e(1_{\mathbb{Z}_{2n}}, a), e_2 = e(a^n, a), \overline{e_1}, \overline{e_2}\}$. Then by Proposition 5.5 we can conclude that $(\mathbb{Z}_{2n}, \mathcal{S})$ is a totally reflected system for any $n \in \mathbb{N}$.

$\diamond$

For a moment, consider the group $\mathbb{Z}_2$. This group has only one nice generating set, namely the one consisting of the single non-identity element of $\mathbb{Z}_2$. The preceding example shows that $\mathbb{Z}_2$ is totally reflected with respect to this generating set. Consequently, $\mathbb{Z}_2$ is totally reflected with respect to *every* one of its nice generating sets. We will see soon that $\mathbb{Z}_2$ is the only nontrivial group having this property.

Lemma 5.8. *Let G be a group generated by $\mathcal{S}$. If $\mathcal{S}$ contains an element which has odd order in the group, then $(G, \mathcal{S})$ is not a totally reflected system.*

Proof. Assume $s \in \mathcal{S}$ has odd order, n, in G. Since $1_G \notin \mathcal{S}$, n must be greater than or equal to 3. We want to show that no color-preserving graph reflection on $\Gamma = \Gamma(G, \mathcal{S})$ can invert the edge $e(1_G, s)$. Suppose, to the contrary, that $L_a\phi \in \mathrm{Aut}_{\text{c.p.}}(\Gamma)$ is a reflection on Γ which inverts the edge $e(1_G, s)$. We showed in Chapter 4 that we must have $a = s$, $\phi(s) = s^{-1}$, and $\phi^2 = \mathrm{id}_G$. Consider the path $\gamma = (e_1, \ldots, e_n)$ in Γ, where $e_i = e(s^{i-1}, s)$ for each $i \in I_n$. Since s has order n in G, $s^0, s^1, s^2, \ldots, s^{n-1}$ are all distinct. If for some integers i and j satisfying $0 \le i < j \le n-1$ we had $s^i = s^j$, then we would also have $1_G = s^{j-i}$. However, the condition on i and j implies that $0 < j - i \le n - 1 - i < n$, contradicting our assumption that n is the order of s. Notice also that $\tau(e_n) = \tau\left(e(s^{n-1}, s)\right) = s^{n-1}s = s^n = 1_G = \iota\left(e(s^0, s)\right) = \iota(e_1)$. Therefore, $\gamma = (e_1, \ldots, e_n)$ is, in fact, a cycle in Γ.

Suppose $e_{k+1} = e(s^k, s)$ is an arbitrary edge in γ, where $0 \le k \le n-1$. In order for $L_s\phi$ to invert e_{k+1}, we must have that $s^k = L_s\phi(s^{k+1}) = ss^{-k-1} = s^{-k}$, implying that $s^{2k} = 1_G$. But this occurs if and only if n divides $2k$. The only value k can take between 0 and $n-1$ so that this is true is $k = 0$, meaning that e_1 is the *only* edge in γ which is inverted by $L_s\phi$. But this is a contradiction! Lemma 4.4 tells us that any cycle in Γ must contain an even number of $L_s\phi$-reflectors. Thus, there can be no color-preserving reflection on Γ which inverts the edge $e(1_G, s)$, and by Proposition 5.5 we can conclude that $(G, \mathcal{S})$ is not a totally reflected system.

□

Lemma 5.9. *Suppose G is a group with generating set $\mathcal{S}$. If $\Gamma = \Gamma(G, \mathcal{S})$ contains a triangle, then $(G, \mathcal{S})$ is not a totally reflected system.*

Proof. Suppose γ is a triangle in Γ and let e be any edge in γ. Then Corollary 4.6 tells us that there does not exist any color-preserving reflection (or any reflection, for that matter) that inverts the edge e. Therefore, $(G, \mathcal{S})$ is not a totally reflected system. □

Corollary 5.10. *Let G be any group with 3 or more elements. Then there exists at least one (nice) generating set $\mathcal{S}$ for G such that $(G, \mathcal{S})$ is not a totally reflected system.*

Proof. Let $\mathcal{S}$ be the generating set for G consisting of every element of G. That is, $\mathcal{S} = G$. Let g_1, g_2, and g_3 be any three distinct elements of G. Then $s_1 = g_1^{-1}g_2, s_2 = g_2^{-1}g_3, s_3 = g_3^{-1}g_1 \in \mathcal{S}$, and so $e_1 = e(g_1, s_1), e_2 = e(g_2, s_2), e_3 = e(g_3, s_3)$ are edges in Γ. Observe the following:

$$\iota(e_1) = g_1 \text{ and } \tau(e_1) = g_1 s_1 = g_1 g_1^{-1} g_2 = g_2$$

$$\iota(e_2) = g_2 \text{ and } \tau(e_2) = g_2 s_2 = g_2 g_2^{-1} g_3 = g_3$$

$$\iota(e_3) = g_3 \text{ and } \tau(e_3) = g_3 s_3 = g_3 g_3^{-1} g_1 = g_1$$

Therefore, $\gamma = (e_1, e_2, e_3)$ is a triangle in Γ, and by Lemma 5.9 we can conclude that $(G, \mathcal{S})$ is not a totally reflected system.

□

Lemmas 5.8 and 5.9 suggest that cycles of odd length in a Cayley graph $\Gamma(G, \mathcal{S})$ make it unlikely that $(G, \mathcal{S})$ will be a totally reflected system. However, a Cayley graph $\Gamma(G, \mathcal{S})$ can contain no odd-length cycles and yet $(G, \mathcal{S})$ can fail to be a totally reflected system. The next example will illustrate this point.

Example 5.11. Consider the quaternion group Q_8 given by the presentation here:

$$Q_8 := \langle -1, i, j, k \mid (-1)^2 = 1 = 1_{Q_8}, i^2 = j^2 = k^2 = ijk = -1 \rangle$$

We can see from the presentation that -1 has order 2 in the group, while i,

j, and k each have order 4. This group consists of eight elements, namely $1, -1, i, -i, j, -j, k$, and $-k$.

Suppose $\mathcal{S}$ is an arbitrary generating set for Q_8. Clearly $\mathcal{S}$ cannot consist of -1 alone, or even -1 together with a single element of order 4, since the square of any order-four element is -1. Therefore, $\mathcal{S}$ must contain two distinct elements of order 4, say a and b. The Cayley graph $\Gamma = \Gamma(Q_8, \mathcal{S})$ must contain a subgraph Γ' whose pictorial representation is depicted in Figure 5.1.

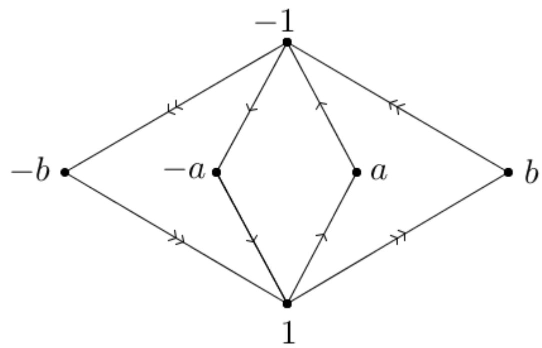

Figure 5.1: Subgraph of $\Gamma(Q_8, \mathcal{S})$

Let $e_1 = e(1, a)$, $e_2 = e(a, a)$, $e_3 = e(-1, a)$, $e_4 = e(-a, a)$, $e_5 = e(1, b)$, and $e_6 = e(b, b)$. Consider the cycles $\gamma_1 = (e_1, e_2, e_3, e_4)$ and $\gamma_2 = (e_5, e_6, e_3, e_4)$. Suppose r is a color-preserving graph reflection on Γ which inverts e_1. Then r must invert another edge in γ_1, since any cycle in Γ must contain an even number of r-reflectors. However, r-reflectors cannot share an endpoint Therefore, the other r-reflector in γ_1 must be e_3. By applying similar reasoning to the cycle γ_2, we can conclude that r must invert e_5 as well. But this is a contradiction! Since e_1 and e_5 share an endpoint, they cannot both be r-reflectors. Consequently, no color-preserving reflection on Γ can invert the edge e_1. Therefore, $(Q_8, \mathcal{S})$ is not a totally reflected system. Since $\mathcal{S}$ was taken to be an arbitrary generating set for Q_8, we can now conclude that the group Q_8 is not totally reflected with respect to any generating set.

◇

The preceding example suggests that there is something about the innate structure of the group Q_8 which makes it difficult (or impossible) to find reflections on its Cayley graphs. This example is the motivation for defining a new term. We will say that a group G is **totally reflected** (or is a **totally reflected group**) if it has a generating set $\mathcal{S}$ for which $(G, \mathcal{S})$ is a totally reflected system. Using this new terminology, we can reinterpret Example 5.11 as saying that Q_8 is not a totally reflected group.

Example 5.12. For any $n \in \mathbb{N}$, let $\mathbb{Z}_{2n+1} = \langle a \mid a^{2n+1} = 1_{\mathbb{Z}_{2n+1}} \rangle$ be the finite cyclic group of order $2n+1$. Let $\mathcal{S}$ be any generating set for $\mathbb{Z}_{2n+1}$. Then from basic group theory we know that $\mathcal{S}$ must contain at least one element of odd order k with $k \geq 3$. By Lemma 5.8, we can conclude that for any generating set $\mathcal{S}$ of $\mathbb{Z}_{2n+1}$, $(\mathbb{Z}_{2n+1}, \mathcal{S})$ is not a totally reflected system. Consequently, for any $n \in \mathbb{N}$, the group $\mathbb{Z}_{2n+1}$ is not totally reflected.

◇

5.2 Direct Products of Totally Reflected Groups

After defining the concept of *totally reflected*, we gave many examples of specific groups which are totally reflected. Our goal now is to discover and understand larger, more general classes of totally reflected groups. The first such class will consist of finite direct products of totally reflected groups.

Let us first establish some conventions that we will use when working with finite direct products. Assume that G and H are groups with identity elements 1_G and 1_H, respectively. Consider the direct product $G \times H$. Traditionally, elements of this direct product are denoted as ordered pairs of the form (g, h), with $g \in G$ and $h \in H$. The group operation is shown here:

$$(g_1, h_1) \cdot (g_2, h_2) = (g_1 g_2, h_1 h_2), \qquad \text{for } g_1, g_2 \in G,\ h_1, h_2 \in H$$

As established previously, the binary operators of G and H are dropped and multiplication within each group is denoted by juxtaposition of elements. For convenience, we will use juxtaposition notation when working with direct products as well. We will write gh in place of (g, h), g in place of $(g, 1_H)$, h in place of $(1_G, h)$, and 1 in place of $(1_G, 1_H)$. The group operation of $G \times H$ can then be summarized as shown here:

$$g_1 h_1 \cdot g_2 h_2 = g_1 h_1 g_2 h_2 = g_1 g_2 h_1 h_2, \qquad \text{for } g_1, g_2 \in G,\ h_1, h_2 \in H$$

The equivalence between internal and external direct products, which is a standard result of abstract algebra, justifies the use of juxtaposition notation when working with direct products.

In a more general finite direct product of the form $G_1 \times G_2 \times \cdots \times G_n$, group elements often will be written in the form $g = g_1 g_2 \cdots g_n$, where $g_i \in G_i$ for $i \in I_n$. If for some $i \in I_n$ we have that $g_i = 1_{G_i}$, we simply will leave out that g_i from the juxtaposed expression of the element g. If $g_i = 1_{G_i}$ for all $i \in \{1, ..., n\}$, then $g = 1$, where 1 is the identity element in the direct product.

One nice consequence of using juxtaposition notation is the ease of working with subsets. If K is a subset of one of the factor groups G_i, then we can also think of K as being a subset of the direct product $G_1 \times G_2 \times \cdots \times G_n$. For example, in a direct product of the form $G \times H$, suppose that K is a subset of H. Then in the standard notation, the set $\{(1_G, k) \mid k \in K\}$ is a subset of $G \times H$. But in our juxtaposed notation, this set would instead be thought of just as $\{k \mid k \in K\}$, or in essence, K itself.

As previously mentioned, the choice of generating set for a group matters when we discuss Cayley graphs and totally reflected systems. If $G_1, G_2, \ldots, G_n$ are groups with generating sets $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_n$, respectively, then one standard generating set for the direct product $G_1 \times G_2 \times \cdots \times G_n$ is formed by taking the union of the generating sets for the individual factors, as seen here:

$$\mathcal{S}_1 \cup \mathcal{S}_2 \cup \cdots \cup \mathcal{S}_n = \bigcup_{i=1}^{n} \mathcal{S}_i$$

Though there are potentially many other generating sets for the finite direct product $G_1 \times G_2 \times \cdots \times G_n$, this standard generating set is the one we typically will use. We will now show that if $(G_1, \mathcal{S}_1), (G_2, \mathcal{S}_2), \ldots, (G_n, \mathcal{S}_n)$ are each totally reflected systems, for $n \geq 2$, then $(G_1 \times G_2 \times \cdots \times G_n, \mathcal{S}_1 \cup \mathcal{S}_2 \cup \cdots \cup \mathcal{S}_n)$ is a totally reflected system. We will first prove this result in the case when $n = 2$. The result for a general finite direct product will then follow easily by induction.

Theorem 5.13. *If G and H are totally reflected groups, then the direct product $G \times H$ is a totally reflected group.*

Proof. Let $\mathcal{S}$ and $\mathcal{T}$ be nice generating sets for G and H, respectively, such that $(G, \mathcal{S})$ and $(H, \mathcal{T})$ are totally reflected systems. We will show that the direct product $G \times H$ is totally reflected by showing that $(G \times H, \mathcal{S} \cup \mathcal{T})$ is a totally reflected system. For convenience, we will abbreviate $\Gamma(G, \mathcal{S})$, $\Gamma(H, \mathcal{T})$, and $\Gamma(G \times H, \mathcal{S} \cup \mathcal{T})$ as $\Gamma(G)$, $\Gamma(H)$, and $\Gamma(G \times H)$, respectively.

Let's begin by defining a *non-rigid* graph morphism $p_G : \Gamma(G \times H) \to \Gamma(G)$. We first will define p_G on the vertex set of $\Gamma(G \times H)$. For any element v in the vertex set of $\Gamma(G \times H)$, v can be expressed in the form gh for some $g \in G$ and some $h \in H$. Then let $p_G(v) = p_G(gh) = g$. We now will define p_G on the edge set of $\Gamma(G \times H)$. For any element e in the edge set of $\Gamma(G \times H)$, e can be expressed either in the form $e(gh, s)$ or in the form $e(gh, t)$, for some $g \in G$, $h \in H$, $s \in \mathcal{S}^{\pm}$, and $t \in \mathcal{T}^{\pm}$. Then let $p_G(e) = e(g, s)$ if $e = e(gh, s)$ and $p_G(e) = g$ if $e = e(gh, t)$. We must check that p_G respects the reversing and initial-vertex functions. To this end, suppose that e is an edge in $\Gamma(G \times H)$. If $e = e(gh, s)$, where $g \in G$, $h \in H$, and $s \in \mathcal{S}^{\pm}$, then $p_G(e) = e(g, s)$. Then $p_G\big(\iota(e)\big) = p_G(gh) = g = \iota\big(p_G(e)\big)$ and $p_G(\overline{e}) = p_G\big(e(ghs, s^{-1})\big) = p_G\big(e(gsh, s^{-1})\big) = e(gs, s^{-1}) = \overline{e(g, s)} = \overline{p_G(e)}$. If $e = e(gh, t)$, where $g \in G$, $h \in H$, and $t \in \mathcal{T}^{\pm}$, then $p_G(e) = g$. Then $p_G\big(\iota(e)\big) = p_G(gh) = g = p_G(e)$ and $p_G(\overline{e}) = p_G\big(e(ght, t^{-1})\big) = g = p_G(e)$. Therefore, p_G does, indeed, respect the reversing and initial-vertex functions, and so p_G is a well-defined non-rigid graph morphism.

Similarly, let's define a non-rigid graph morphism $p_H : \Gamma(G \times H) \to \Gamma(H)$. We first will define p_H on the vertex set of $\Gamma(G \times H)$. Again, for any element v in the vertex set of $\Gamma(G \times H)$, v can be expressed in the form gh for some $g \in G$ and some $h \in H$. Then let $p_H(v) = p_H(gh) = h$. We now will define p_H on the edge set of $\Gamma(G \times H)$. Again, for any element e in the edge set of $\Gamma(G \times H)$, e can be expressed either in the form $e(gh, s)$ or in the form $e(gh, t)$, for some $g \in G$, $h \in H$, $s \in \mathcal{S}^{\pm}$, and $t \in \mathcal{T}^{\pm}$. Then let $p_H(e) = h$ if $e = e(gh, s)$ and $p_H(e) = e(h, t)$ if $e = e(gh, t)$. An argument similar to the one above would give us that p_H respects the reversing and initial-vertex functions. Therefore, p_H is a well-defined non-rigid graph morphism.

Let $s \in \mathcal{S}$ be arbitrary. Consider the edge $e(1, s) = e(1_G 1_H, s)$ in $\Gamma(G \times H)$. Then $p_G(e(1, s)) = e(1_G, s)$ is an edge in $\Gamma(G)$. Since $(G, \mathcal{S})$ is a totally reflected system, there is a color-preserving graph reflection of the form $L_s\alpha$, with $\alpha \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$ satisfying $\alpha(s) = s^{-1}$ and $\alpha^2 = \mathrm{id}_G$, which inverts the edge $e(1_G, s)$ in $\Gamma(G)$.

Now, α is a bijective function on G, but we can extend it to a function α^* on all of $G \times H$ by defining $\alpha^*(1) = 1$ and $\alpha^*(gh) = \alpha(g)h$ for all $g \in G$ and $h \in H$. Suppose $gh \in G \times H$ is arbitrary, where $g \in G$ and $h \in H$. Since α is bijective, $\alpha^{-1}(g)$ is well defined. Then observe that $\alpha^*\left(\alpha^{-1}(g)h\right) = \alpha\left(\alpha^{-1}(g)\right)h = gh$. Therefore, α^* is surjective. Also, $\alpha(gh) = \alpha(g)h = 1$ if and only if $\alpha(g) = 1_G$ and $h = 1_H$. Since α is injective, $\alpha(g) = 1_G$ occurs if and only if $g = 1_G$, and so $gh = 1_G 1_H = 1$. Therefore, $\alpha^*(gh) = 1$ occurs if and only if $gh = 1$, and so α^* is injective. Suppose that $g, g' \in G$ and $h, h' \in H$. Observe the following:

$$\begin{aligned} \alpha^*(ghg'h') = \alpha^*(gg'hh') = \alpha(gg')hh' = \alpha(g)\alpha(g')hh' \\ = \alpha(g)h\alpha(g')h' = \alpha^*(gh)\alpha^*(g'h') \end{aligned}$$

Therefore, we can now conclude that α^* is an automorphism on $G \times H$. Furthermore, α^* fixes the set $\mathcal{S}^{\pm} \cup \mathcal{T}^{\pm}$ since α fixes the set $\mathcal{S}^{\pm}$ and since the extension α^* is defined so as to fix the set $\mathcal{T}^{\pm}$. Therefore, $\alpha^* \in \mathrm{Aut}(G \times$

$H, \mathcal{S}^{\pm} \cup \mathcal{T}^{\pm})$, and by Theorem 3.8 we know that $L_s\alpha^*$ is a color-preserving graph automorphism acting on $\Gamma(G \times H)$.

We can see that $L_s\alpha^*$ inverts the edge $e(1,s)$ in $\Gamma(G\times H)$, since $L_s\alpha^*(1) = s$ and $L_s\alpha^*(s) = s\alpha(s) = ss^{-1} = 1$. Our next goal is to show that $L_s\alpha^*$ is a reflection on $\Gamma(G \times H)$.

First, we must show that $L_s\alpha^*$ has order 2. For any vertex $gh \in \Gamma(G\times H)$, notice the following:

$$\begin{aligned}
(L_s\alpha^*)^2(gh) &= L_s\alpha^* L_s\alpha^*(gh) \\
&= s\alpha^*\big(s\alpha^*(gh)\big) \\
&= s\alpha^*\big(s\alpha(g)h\big) \\
&= s\alpha\big(s\alpha(g)\big)h \quad \text{(since } s\alpha(g) \in G\text{)} \\
&= \big(L_s\alpha\big)^2(g)h \\
&= gh \quad \text{(since } (L_s\alpha)^2 = \mathrm{id}_{\Gamma(G)}\text{)}
\end{aligned}$$

Thus, $(L_s\alpha^*)^2 = \mathrm{id}_{\Gamma(G\times H)}$ and so $L_s\alpha^*$ has order 2.

Second, we need to identify those edges in $\Gamma(G \times H)$ which are inverted by $L_s\alpha^*$. To this end, let $\mathcal{E} = \{$edges in $\Gamma(G)$ which are inverted by $L_s\alpha\} = \Gamma(G)_{L_s\alpha}$. Let $\widetilde{\mathcal{E}} = \{$edges, e, in $\Gamma(G \times H)$ such that $p_G(e) \in \mathcal{E}\}$. Since $e(1_G, s)$ is inverted by $L_s\alpha$ acting on $\Gamma(G)$, $e(1_G, s) \in \mathcal{E}$. Also, $e(1,s) \in \widetilde{\mathcal{E}}$ since $p_G(e(1,s)) = e(1_G,s) \in \mathcal{E}$.

Note that in order for an edge $e \in \widetilde{\mathcal{E}}$ to satisfy $p_G(e) \in \mathcal{E}$, we must have that $c(e) \in \mathcal{S}^{\pm}$. We will now show that each edge in $\widetilde{\mathcal{E}}$ is inverted by $L_s\alpha^*$. Consider an arbitrary element of $\widetilde{\mathcal{E}}$, say $e(gh, s')$, where $g \in G$, $h \in H$, and $s' \in \mathcal{S}^{\pm}$. Then $p_G\big(e(gh,s')\big) = e(g,s') \in \mathcal{E}$. Since $e(g,s')$ is inverted by $L_s\alpha$, we must have that $s\alpha(g) = gs'$ and $\alpha(s') = (s')^{-1}$. Then, notice the following:

$$\begin{aligned} L_s\alpha^*\big(e(gh, s')\big) &= e\big(s\alpha^*(gh), \alpha^*(s')\big) \\ &= e\big(s\alpha(g)h, \alpha(s')\big) \\ &= e\big(gs'h, (s')^{-1}\big) \\ &= e\big(ghs', (s')^{-1}\big) \\ &= \overline{e(gh, s')} \end{aligned}$$

Since the edge $e(gh, s')$ was chosen arbitrarily from $\widetilde{\mathcal{E}}$, we have now shown that each edge in $\widetilde{\mathcal{E}}$ is inverted by $L_s\alpha^*$.

Finally, we must show that the edges of $\widetilde{\mathcal{E}}$, if removed, separate the graph $\Gamma(G \times H)$. Suppose that v_1 and v_2 are vertices in $\Gamma(G \times H)$. Then $v_1 = g_1h_1$ and $v_2 = g_2h_2$ for some $g_1, g_2 \in G$ and some $h_1, h_2 \in H$. Any path in $\Gamma(G \times H)$ from v_1 to v_2 projects down, by way of p_G, to a path (perhaps of length 0) in $\Gamma(G)$ from g_1 to g_2. We showed previously that $e(1, s) \in \widetilde{\mathcal{E}}$. Let γ be any path from 1 to s in $\Gamma(G \times H)$. We claim that γ must contain at least one edge which is in $\widetilde{\mathcal{E}}$.

To prove the claim, suppose to the contrary that γ does not contain an edge in $\widetilde{\mathcal{E}}$. Then $p_G(\gamma)$ is a path in $\Gamma(G)$ from 1_G to s which contains no edges from $\mathcal{E}$. Note that $p_G(\gamma)$ is not an empty path, since $1_G \neq s$. Therefore, $p_G(\gamma)$ is a (nonempty) path in $\Gamma(G)$ from 1_G to s which contains no edges which are inverted by $L_s\alpha$, implying that 1_G and s are in the same connected component of $\Gamma(G) - \mathcal{E}$. But this is a contradiction! The vertices 1_G and s must be in different connected components of $\Gamma(G) - \mathcal{E}$ since $e(1_G, s) \in \mathcal{E}$. Thus, any path in $\Gamma(G \times H)$ from 1 to s must contain an edge which is in $\widetilde{\mathcal{E}}$.

Consequently, $\widetilde{\mathcal{E}}$ separates $\Gamma(G \times H)$.

Therefore, $\widetilde{\mathcal{E}}$ is a set of edges in $\Gamma(G \times H)$, each of which is inverted by $L_s\alpha^*$, and $\Gamma(G \times H) - \widetilde{\mathcal{E}}$ is disconnected. So $L_k\alpha^*$ is indeed a color-preserving reflection on $\Gamma(G \times H)$ and it inverts the specified edge $e(1, s)$. Since $s \in \mathcal{S}$ was arbitrary, this implies that for any $s \in \mathcal{S}$, there exists a color-preserving reflection on $\Gamma(G \times H)$ inverting the edge $e(1, s)$. A symmetric argument would show that there exists a color-preserving reflection on $\Gamma(G \times H)$ inverting any edge of the form $e(1, t)$, where $t \in \mathcal{T}$. Therefore, for any $u \in \mathcal{S}^{\pm} \cup \mathcal{T}^{\pm}$, there exists a color-preserving reflection on $\Gamma(G \times H)$ which inverts the edge $e(1, u)$. By Proposition 5.5, we can now conclude that $(G \times H, \mathcal{S}^{\pm} \cup \mathcal{T}^{\pm})$ is a totally reflected system, implying that $G \times H$ is a totally reflected group.

□

Corollary 5.14. *If $G_1, G_2, ..., G_n$ are totally reflected groups for $n \geq 2$, then the direct product $G_1 \times G_2 \times \cdots \times G_n$ is a totally reflected group.*

Example 5.15. Consider the group $\mathbb{Z} \times \mathbb{Z} = \langle a, b \mid ab = ba \rangle$. Figure 5.2 shows the Cayley graph $\Gamma = \Gamma(\mathbb{Z} \times \mathbb{Z}, \{a, b\})$.

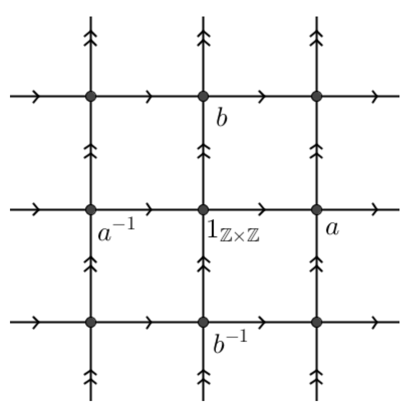

Figure 5.2: A portion of the Cayley graph $\Gamma(\mathbb{Z} \times \mathbb{Z}, \{a, b, \})$

The preceding theorem tells us that $(\mathbb{Z}\times\mathbb{Z},\{a,b\})$ must be totally reflected, since $(\mathbb{Z},\{a\})$ and $(\mathbb{Z},\{b\})$ are totally reflected, as we showed in Example 5.4. Define a group homomorphism ϕ_a on $\mathbb{Z}\times\mathbb{Z}$ by stating that $\phi_a(a) = a^{-1}$ and $\phi_a(b) = b$. Likewise, define a group homomorphism ϕ_b on $\mathbb{Z}\times\mathbb{Z}$ by stating that $\phi_b(a) = a$ and $\phi_b(b) = b^{-1}$. The color-preserving reflections on Γ given by $L_a\phi_a$, $L_{a^{-1}}\phi_a$, $L_b\phi_b$, and $L_{b^{-1}}\phi_b$ invert the edges from $1_{\mathbb{Z}\times\mathbb{Z}}$ to a, a^{-1}, b, and b^{-1}, respectively. Therefore, for each edge e incident to the identity vertex in Γ, there exists a color-preserving reflection on Γ which inverts e. Then by Proposition 5.5 we can confirm that $(\mathbb{Z}\times\mathbb{Z},\{a,b\})$ is, in fact, totally reflected.

$\diamond$

5.3 Finitely Generated Abelian Groups

From the preceding section, we know that any finite direct product of totally reflected groups is totally reflected. Indeed, Theorem 5.13 proved the following implication:

$$G_1, G_2, \ldots, G_n \text{ totally reflected } \Longrightarrow G_1 \times G_2 \times \cdots \times G_n \text{ totally reflected}$$

However, the reverse implication is generally not true. To see this, recall that by Example 5.7 we know that $\mathbb{Z}_2 \times \mathbb{Z}_3$ is totally reflected since it is isomorphic to $\mathbb{Z}_6$. But Example 5.12 shows us that $\mathbb{Z}_3$ itself is not totally reflected. Therefore, a finite direct product of groups being totally reflected does not necessarily imply that the individual factor groups are themselves totally reflected.

Finitely-generated abelian groups will provide us with a setting in which the reverse implication of Theorem 5.13 is sometimes true. Let us begin by recalling an important result from algebra.

Theorem 5.16 (Fundamental Theorem of Finitely Generated Abelian Groups)**.** *Let G be a finitely generated abelian group. Then G can be expressed as a finite direct product of cyclic groups as shown here:*

$$G \cong \mathbb{Z}_{n_1} \times \mathbb{Z}_{n_2} \times \cdots \times \mathbb{Z}_{n_k} \times \mathbb{Z}^r$$

where k and r are nonnegative integers, $n_i \geq 2$ for each $i \in I_k$, and $n_i | n_{i+1}$ for each $i \in I_{k-1}$ in the event that $k \geq 2$.

Furthermore, if $G \cong \mathbb{Z}_{m_1} \times \mathbb{Z}_{m_2} \times \cdots \mathbb{Z}_{m_l} \times \mathbb{Z}^q$, where $m_1, ..., m_l, l$, and q satisfy the same conditions which were satisfied by $n_1, ..., n_k, k$, and r, then $k = l$, $n_i = m_i$ for all $i \in I_k$, and $r = q$.

The direct-product decomposition of a finitely generated abelian group such as that given in the above theorem is sometimes referred to as the *invariant-factor decomposition* for G. The values $n_1, n_2, ..., n_k$ are called the *invariant factors* of G, with n_1 being the *first invariant factor* of G. The value r is called the *rank* of G. When $r = 0$, we interpret this to mean that $G \cong \mathbb{Z}_{n_1} \times \cdots \times \mathbb{Z}_{n_k}$. When $k = 0$ in this decomposition, we interpret this to mean that $G \cong \mathbb{Z}^r$. In this case, we will say that the first invariant factor of G is 0. Throughout this section, we will assume that G is a nontrivial group. Consequently, k and r cannot both be zero.

We will soon prove that a finitely generated abelian group is totally reflected if and only if each factor group in its invariant-factor decomposition is totally reflected. First we must prove the following lemma. For the remainder of this section, we will assume that we are in the setting established in Theorem 5.16. Specifically, if G is a finitely generated abelian group, then we will assume that it has an invariant-factor decomposition such as that in Theorem 5.16.

Lemma 5.17. *Let G be a finitely generated abelian group. Every generating set for G must have at least $k+r$ elements, where k is the number of invariant factors of G and r is the rank of G.*

Proof. Suppose that $\mathcal{S} = \{s_1, ..., s_j\} \neq \varnothing$ is any nice generating set for G. We want to show that it must be the case that $j \geq k + r$. Let $\phi : G \to \mathbb{Z}_{n_1} \times \cdots \times \mathbb{Z}_{n_k} \times \mathbb{Z}^r$ be an isomorphism. For each $i \in I_j$, let $t_i = \phi(s_i)$. Then the set $\mathcal{T} = \{t_1, \ldots, t_j\}$ is a generating set for $\mathbb{Z}_{n_1} \times \cdots \times \mathbb{Z}_{n_k} \times \mathbb{Z}^r$, since ϕ is an isomorphism.

Suppose for a moment that $k > 0$. Let p be a prime number which divides n_1. In the event that $k \geq 2$, p divides not only n_1 but also n_i for each $i \in I_k$, since $n_i | n_{i+1}$ for each $i \in I_{k-1}$. Let $d_i = n_i/p \in \mathbb{Z}$ for each $i \in I_k$. The group $\mathbb{Z}_{n_i}$ must contain subgroups which are isomorphic to $\mathbb{Z}_p$ and $\mathbb{Z}_{d_i}$ (and which we denote, without loss of clarity, as $\mathbb{Z}_p$ and $\mathbb{Z}_{d_i}$, respectively). Furthermore, the subgroup $\mathbb{Z}_{d_i}$ is normal in $\mathbb{Z}_{n_i}$, since $\mathbb{Z}_{n_i}$ is abelian. So the quotient $\mathbb{Z}_{n_i}/\mathbb{Z}_{d_i}$ is a group, and this quotient group is isomorphic to $\mathbb{Z}_p$. Therefore, for each $i \in I_k$, there exists a surjective homomorphism from $\mathbb{Z}_{n_i}$ onto $\mathbb{Z}_p$, which can be achieved by applying the standard quotient homomorphism, followed by an isomorphism.

There also exists a surjective homomorphism from $\mathbb{Z}$ onto $\mathbb{Z}_p$, say ϕ, defined by stating that $\phi(a) = a(\text{mod } p)$ for any integer a. Here, p can be any prime number greater than or equal to 2. We will use the same p as above if $k > 0$, but if $k = 0$ we can simply choose p to be any prime number greater than or equal to 2.

We have shown that there exists a prime number $p \geq 2$ such that each factor group in the invariant factor decomposition for G can be mapped via a surjective homomorphism to $\mathbb{Z}_p$. Therefore, there also exists a surjective homomorphism $\gamma : \mathbb{Z}_{n_1} \times ... \times \mathbb{Z}_{n_k} \times \mathbb{Z}^r \to (\mathbb{Z}_p)^{k+r}$. Since p is a prime, $\mathbb{Z}_p$ is a field, meaning $(\mathbb{Z}_p)^{k+r}$ can be viewed as a vector space of dimension $k+r$ over the field $\mathbb{Z}_p$. Since $\mathcal{T} = \{t_1, ..., t_j\}$ is a generating set for $\mathbb{Z}_{n_1} \times \cdots \times \mathbb{Z}_{n_k} \times \mathbb{Z}^r$, the set $\{\gamma(t_1), ..., \gamma(t_j)\}$ is a generating set for $(\mathbb{Z}_p)^{k+r}$. In other words, $\{\gamma(t_1), ..., \gamma(t_j)\}$ is a spanning set for the $(k+r)$-dimensional vector space $(\mathbb{Z}_p)^{k+r}$. Therefore, the spanning set $\{\gamma(t_1), ..., \gamma(t_j)\}$ must contain at least $k+r$ distinct elements, implying that $j \geq k+r$ and completing the proof.

□

Theorem 5.18. *Let G be a finitely generated abelian group. Then G is totally reflected if and only if its first invariant factor is even.*

Proof. ($\Longleftarrow$) Assume that the first invariant factor of G is even. Therefore, $n_1 \geq 2$ must be even in the event that $k > 0$, implying that $n_i \geq 2$ is even for each $i \in I_k$. We have shown previously that each finite cyclic group of even order and the infinite cyclic group are totally reflected. Therefore, we can conclude by Corollary 5.14 that $G \cong \mathbb{Z}_{n_1} \times \cdots \times \mathbb{Z}_{n_k} \times \mathbb{Z}^r$ is totally reflected.

($\Longrightarrow$) Assume G is totally reflected, and suppose that $\mathcal{S} = \{s_1, ..., s_j\}$ is a nice generating set for G such that $(G, \mathcal{S})$ is a totally reflected system. By Lemma 5.17, we know that $j \geq k+r$. If $k = 0$, then $G \cong \mathbb{Z}^r$ has first invariant factor of 0, which is even. Therefore, assume for the remainder of the proof that $k > 0$. We will show that n_1 must be even. To the contrary, suppose that n_1 is odd.

First, suppose that $\mathcal{S}$ contains only one element. Since we are assuming that $k > 0$, G is not an infinite cyclic group. Therefore, the single element of $\mathcal{S}$ must be a torsion element, which implies that G is a finite cyclic group and that r equals 0. If $k \geq 2$, then n_1 divides n_i for each $i \in I_k$. Then for any $i, j \in I_k$, $\gcd(n_i, n_j) \geq n_1 > 1$. Therefore, $\mathbb{Z}_{n_1} \times \cdots \times \mathbb{Z}_{n_k}$ is cyclic only when $k = 1$, which now gives us that G must be isomorphic to $\mathbb{Z}_{n_1}$. But since n_1 is odd, $\mathbb{Z}_{n_1}$, and thus G, cannot be totally reflected, which contradicts our fundamental assumption. Therefore, n_1 must, in fact, be even, which proves the lemma in the case where $\mathcal{S}$ contains only one element.

For the remainder of the proof, we will assume that $\mathcal{S}$ contains more than one element, meaning that $j \geq 2$. By Lemma 5.8, we know that $\mathcal{S}$ cannot contain any elements of odd order in G. Therefore, for each $i \in I_j$ we must have that s_i has either infinite order or even order of 2 or greater in the group. For each $i \in I_j$, let $H_{s_i} \coloneqq \langle \mathcal{S} - \{s_i\} \rangle$. Since $\mathcal{S}$ is a nice generating set with more than 1 element, the set H_{s_i} must be a nontrivial subgroup of G.

We claim that there exists at least one $i \in I_j$ such that $\langle s_i \rangle \cap H_{s_i} \neq \{1_G\}$. To prove the claim, suppose that instead $\langle s_i \rangle \cap H_{s_i} = \{1_G\}$ for all $i \in I_j$. For any $i, i' \in I_j$, $i \neq i'$, we have that $s_{i'} \in H_{s_i}$, which implies that $\langle s_i \rangle \cap \langle s_{i'} \rangle \subseteq \langle s_i \rangle \cap H_{s_i} = \{1_G\}$. Therefore, G can be expressed as an (internal) direct product:

$$G \cong \langle s_1 \rangle \times \cdots \times \langle s_j \rangle$$

Now, suppose that $s_i \in \mathcal{S}$ is a generator with finite order. Then, as established before, s_i must have even order of 2 or greater. Therefore, $\langle s_i \rangle$

contains a subgroup, say D, which is isomorphic to $\mathbb{Z}_{d_i}$, where $d_i = \frac{|s_i|}{2}$. Furthermore, D is normal in $\langle s_i \rangle$, since this group is abelian. Then there exists a surjective homomorphism from $\langle s_i \rangle$ onto the quotient group $\langle s_i \rangle / D$, namely the standard quotient homomorphism. But $\langle s_i \rangle / D \cong \mathbb{Z}_2$, and so there exists a surjective homomorphism from $\langle s_i \rangle$ onto $\mathbb{Z}_2$. Such a homomorphism exists for any $s_i \in \mathcal{S}$ which has finite order. Let us now consider a generator $s_i \in \mathcal{S}$ which has infinite order. Then $\langle s_i \rangle \cong \mathbb{Z}$, and there exists a surjective homomorphism from this infinite cyclic group onto $\mathbb{Z}_2$, namely the map defined by sending $(s_i)^n$ to $n (\text{mod } 2)$. Such a homomorphism exists for any $s_i \in \mathcal{S}$ which has infinite order. Therefore, we have shown that for *any* $s_i \in \mathcal{S}$, there exists a surjective homomorphism from $\langle s_i \rangle$ onto $\mathbb{Z}_2$. Consequently, there exists a surjective homomorphism $\delta : G \to (\mathbb{Z}_2)^j$, since $G \cong \langle s_1 \rangle \times \cdots \times \langle s_j \rangle$. However, recall that we also have that $G \cong \mathbb{Z}_{n_1} \times \cdots \times \mathbb{Z}_{n_k} \times \mathbb{Z}^r$ and that n_1 is odd. Since no surjective homomorphism exists from $\mathbb{Z}_{n_1}$ onto $\mathbb{Z}_2$, we must have that $j \leq k - 1 + r$, since δ maps G surjectively onto $(\mathbb{Z}_2)^j$. But we've now shown that $j \geq k + r$ and $j \leq k - 1 + r$, which is impossible. Therefore, our claim is true: there must exist at least one $i \in I_j$ such that $\langle s_i \rangle \cap H_{s_i} \neq \{1_G\}$.

Suppose $s \in \mathcal{S}$ is one such generator satisfying $\langle s \rangle \cap H_s \neq \{1_G\}$. Then there exists a smallest positive integer power of s which is a non-identity element of H_s, say $s^m \neq 1_G$, where $m \geq 1$. That is, if l is an integer with $0 < l < m$, then $s^l \notin H_s$. In particular, $s^l \neq 1_G$ for any integer l with $0 < l < m$. We will show that the index of the subgroup H_s inside of G is m. Because G is abelian and because of the way H_s is defined, any element of G can be written in the form $s^i g$ where $g \in H_s$ and where $0 \leq i < |s|$. Therefore, the cosets of H_s in G are all of the form $s^i H_s$ where $0 \leq i < |s|$.

If $m = 1$, then by definition of m we have that $s \in H_s$ and $s \neq 1$. Therefore, H_s is the only coset of H_s in G. If $m = 2$, we have that $s^2 \in H_s$ and $s^2 \neq 1$. Then H_s and sH_s are the only cosets of H_s in G, which are of course distinct since $s \notin H_s$. Suppose now that $m \geq 3$. Then $s^i H_s$ are the cosets of H_s in G for $i \in \overline{I}_{m-1}$, and all m of these cosets are distinct. To see that they are distinct, suppose that $i, j \in \overline{I}_{m-1}$, $i < j$. If $s^i H_s = s^j H_s$, we would have that $s^{-i} s^j = s^{j-i} \in H_s$. But $0 \leq j - i \leq m - 1 < m$, since i and j are both integers between 0 and $m - 1$, with $i < j$. But then this contradicts m being the smallest nontrivial power of s which is an element of H_s. Therefore, no matter what m is, there will be precisely m distinct cosets of H_s in G which are of the form $s^i H_s$ for $i \in \overline{I}_{m-1}$.

By definition of H_s, s^m can be written as a word $t_1 t_2 \cdots t_p$, where each $t_i \in \mathcal{S}^{\pm} - \{s, s^{-1}\}$. We may assume that for any $i, j \in I_p$ with $i < j$, $t_{i+1} t_{i+2} \cdots t_{j-1} t_j \neq 1_G$, since if this did occur for some $i, j \in I_p$ with $i < j$, we could write s^m as $t_1 \cdots t_i t_{j+1} \cdots t_p$ instead. Since $s^m \neq 1$ and since none of the t_i equals s or s^{-1}, we know that $p \geq 2$. Let us now define edges e_i in $\Gamma(G, \mathcal{S})$ for $i \in \overline{I}_{p+m-1}$. For $i \in \overline{I}_{m-1}$, define $e_i := e(s^i, s)$. For $i \in \{m, m+1, \ldots, p+m-1\}$ (which must contain at least two elements, since $p \geq 2$ and $m \geq 1$), define $e_i := e(t_1 \cdots t_{p+m-i}, (t_{p+m-i})^{-1})$. Note the following relationships. For $i \in \{m, m+1, ..., p+m-2\}$, $\tau(e_i) = t_1 \cdots t_{p+m-i} \cdot (t_{p+m-i})^{-1} = t_1 \cdots t_{p+m-(i+1)} = \iota(e_{i+1})$. If $i = m-1$, $\tau(e_i) = \tau(e_{m-1}) = s^{m-1} s = s^m = t_1 t_2 \cdots t_p = \iota(e_m) = \iota(e_{i+1})$. If $m \geq 2$ and $i \in \overline{I}_{m-2}$, $\tau(e_i) = s^{i+1} = \iota(e_{i+1})$. Thus, $e_0 e_1 \cdots e_{p+m-1}$ is a well-defined path in $\Gamma(G, \mathcal{S})$. Furthermore, $\iota(e_1) = s^0 = 1$ and $\tau(e_{p+m-1}) = 1$, since $p + m - (p + m - 1) = 1$. Note that $s^0 = 1_G, s^1, \ldots, s^{m-1}, t_1, t_1 t_2, \ldots, t_1 t_2 \cdots t_p$ are all distinct elements of G. If $s^i = s^j$ for some $i, j \in \overline{I}_{m-1}$ with $i < j$, then we would have $0 < j - i \leq m - 1 < m$ and $s^{j-i} = 1_G$, which we previously showed could not happen. If $t_1 t_2 \cdots t_i = t_1 t_2 \cdots t_j$ for some $i, j \in I_p$ with $i < j$, then we would have $t_{i+1} t_{i+2} \cdots t_{j-1} t_j = 1_G$, which we previously showed could not happen. Lastly, for any $i \in \overline{I}_{m-1}$ and for any $j \in I_p$, we cannot have $s^i = t_1 t_2 \cdots t_j$, since the cosets $s^i H_s$ and H_s are distinct, as previously shown. Therefore, for any $i, j \in \overline{I}_{p+m-1}$, $i \neq j$, we have $\iota(e_i) \neq \iota(e_j)$. Therefore, we can now conclude that $e_0 e_1 \cdots e_{p+m-1}$ is a cycle in $\Gamma(G, \mathcal{S})$.

Since $(G, \mathcal{S})$ is totally reflected, there exists a color-preserving reflection of the form $L_s \alpha$ acting on $\Gamma(G, \mathcal{S})$, with $\alpha \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$ satisfying $\alpha(s) = s^{-1}$ and $\alpha^2 = \mathrm{id}_G$, which inverts the edge $e_0 = e(1_G, s)$ in $\Gamma(G, \mathcal{S})$. Furthermore, since the group G is abelian, Lemma 5.3 tells us that α must fix all the elements of the set $\mathcal{S}^{\pm} - \{s, s^{-1}\}$. Consequently, for any element $g \in H_s$, $\alpha(g) = g$. Consider an arbitrary edge of the form $e(g, t)$ in $\Gamma(G, \mathcal{S})$, with $g \in H_s$ and $t \in \mathcal{S}^{\pm} - \{s, s^{-1}\}$. Then

$$L_s \alpha \left(e(g, t) \right) = e \left(s\alpha(g), \alpha(t) \right) = e(sg, t)$$

In order for the edge $e(g, t)$ to be inverted by the reflection $L_s \alpha$, we would need $gt = sg$, which would imply (since G is abelian) that $t = s$. But we were assuming that $t \neq s$ and $t \neq s^{-1}$. Therefore, the edge $e(g, t)$ cannot be inverted by $L_s \alpha$. In particular, the edges $e_m, e_{m+1}, \ldots, e_{p+m-1}$ as defined previously cannot be inverted by $L_s \alpha$.

Now consider an edge $e_i = e(s^i, s)$ in $\Gamma(G, \mathcal{S})$, with $i \in \overline{I}_{m-1}$. Then

$$L_s \alpha \left(e(s^i, s) \right) = e \left(s\alpha(s^i), \alpha(s) \right) = e \left(s s^{-i}, s^{-1} \right) = e \left(s^{1-i}, s^{-1} \right)$$

In order for the edge $e(s^i, s)$ to be inverted by the reflection $L_s\alpha$, we would need $s^i = \iota(e(s^i, s)) = \tau(e(s^{1-i}, s^{-1})) = s^{-i}$, or in other words $s^{2i} = 1_G$. This certainly occurs when $i = 0$, though we already are assuming that the edge $e(1_G, s)$ is inverted by $L_s\alpha$. If $m \geq 2$, the set $\{e_0, ..., e_{m-1}\}$ contains more than one edge, but we claim that e_0 is the only edge in the set which is inverted by $L_s\alpha$.

To verify this claim, first recall that we are assuming that m is the smallest positive integer power of s which is an element of H_s. Now, suppose $m \geq 2$. We want to show that none of the edges $e_1, ..., e_{m-1}$ is inverted by $L_s\alpha$. For e_i to be inverted when $i \in I_{m-1}$, recall that we need $s^{2i} = 1_G$. If $i < m/2$, then $2i < m$, and we know that $s^{2i} \neq 1_G$. So, e_i is not inverted if $i < m/2$. If $i = m/2$ (which only happens if m is even), then $2i = m$, but specifically $s^m \neq 1_G$, by definition of m. So e_i is not inverted if $i = m/2$. Now suppose that $i > m/2$. Then we have that $m/2 < i \leq m-1$, which implies that $m < 2i \leq 2m-2$, and so $0 < 2i - m \leq m-2$. Then $s^{2i} = s^{2i-m}s^m \in s^{2i-m}H_s$, and $s^{2i-m}H_s \neq H_s$, since $0 < 2i - m \leq m - 2$. Thus, $s^{2i} \neq 1_G$, since $1_G \notin s^{2i-m}H_s$. So, e_i is not inverted if $i > m/2$. Therefore, if $m > 1$, none of $e_1, ..., e_{m-1}$ is inverted.

We have now shown that e_0 is the only edge in the cycle $e_0e_1\cdots e_{p+m-1}$ which is inverted by $L_s\alpha$, meaning that $L_s\alpha$ does not separate the graph $\Gamma(G, \mathcal{S})$. But this contradicts our assumption that $L_s\alpha$ is a reflection on $\Gamma(G, \mathcal{S})$. Therefore, tracing our argument backwards, we see that n_1 cannot, in fact, be odd. Therefore, n_1 must be even, which completes this direction of our overall proof.

Therefore, G is totally reflected if and only if its first invariant factor is even. □

5.4 Free Products of Totally Reflected Groups

The third general class of totally reflected groups which we will explore consists of finite free products of totally reflected groups. We will assume that the reader is somewhat familiar with the free product construction. However, when working with finite free products of groups, we will make use of some of the conventions and terminology set forth by Magnus, Karrass, and Solitar [6]. Though there are different ways of defining a free product of groups, we will use the group-presentation definition.

Assume that G and H are groups with identity elements 1_G and 1_H, respectively. Also, assume G and H have presentations $\langle \mathcal{S} \mid \mathcal{R}_G \rangle$ and $\langle \mathcal{T} \mid \mathcal{R}_H \rangle$, respectively, where $\mathcal{S}$ and $\mathcal{T}$ are nice generating sets. Define the free product of G and H to be the group denoted by $G * H$ with presentation $\langle \mathcal{S} \cup \mathcal{T} \mid \mathcal{R}_G \cup \mathcal{R}_H \rangle$. The identity element of $G * H$, usually called the **empty word**, will be denoted by 1. In the free product $G * H$, the elements 1_G and 1_H are identified with the empty word. For example, if $s \in \mathcal{S}$, then we will write $ss^{-1} = 1$ instead of $ss^{-1} = 1_G$.

Let W be a word in $\mathcal{S}^{\pm} \cup \mathcal{T}^{\pm}$, as shown here:

$$W = x_1 x_2 \cdots x_n, \qquad \text{where } x_i \in \mathcal{S}^{\pm} \cup \mathcal{T}^{\pm} \text{ for all } i \in I_n.$$

An **$\mathcal{S}$-syllable** of W is any subword of the form $x_i x_{i+1} \cdots x_{i+n_i}$, where $i \in I_n$, $n_i \in \overline{I}_{n-i}$, the letters $x_i, x_{i+1}, ..., x_{i+n_i}$ are all elements of $\mathcal{S}^{\pm}$, x_{i-1} is an element of $\mathcal{T}^{\pm}$ in the event that $i \geq 2$, and x_{i+n_i+1} is an element of $\mathcal{T}^{\pm}$ in the event that $i + n_i \leq n - 1$. Similarly, a **$\mathcal{T}$-syllable** of W is any subword of the form $x_i x_{i+1} \cdots x_{i+n_i}$, where $i \in I_n$, $n_i \in \overline{I}_{n-i}$, the letters $x_i, x_{i+1}, ..., x_{i+n_i}$ are all elements of $\mathcal{T}^{\pm}$, x_{i-1} is an element of $\mathcal{S}^{\pm}$ in the event that $i \geq 2$, and x_{i+n_i+1} is an element of $\mathcal{S}^{\pm}$ in the event that $i + n_i \leq n - 1$.

As we move from left to right in the expression $x_1 x_2 \cdots x_n$ we can derive a unique sequence of alternating $\mathcal{S}$-syllables and $\mathcal{T}$-syllables $(w_1, w_2, ..., w_r)$. This sequence will be called the **syllable decomposition** of the word W. The positive integer r will be called the **syllable length** of W and will be denoted as $l(W)$.

An example will help us to illustrate these concepts. Suppose that $\mathcal{S} = \{s_1, s_2\}$ and $\mathcal{T} = \{t_1, t_2, t_3\}$. Consider the following word in $\mathcal{S}^{\pm} \cup \mathcal{T}^{\pm}$:

$$W = s_1(s_2)^{-1}s_1t_1t_1t_1(s_1)^{-1}(s_1)^{-1}t_3t_2 = s_1(s_2)^{-1}s_1(t_1)^3(s_1)^{-2}t_3t_2$$

The syllable decomposition for W is (w_1, w_2, w_3, w_4), where $w_1 = s_1(s_2)^{-1}s_1$, $w_2 = (t_1)^3$, $w_3 = (s_1)^{-2}$ and $w_4 = t_3t_2$. Therefore, $l(W) = 4$.

We can generalize the construction and terminology involved in the free product of two groups to the case of the free product of any finite number of groups. Assume $G_1, \ldots, G_n$ are groups. Assume that for each $i \in I_n$ we have $G_i = \langle \mathcal{S}_i \mid \mathcal{R}_i \rangle$, where $\mathcal{S}_i$ is a nice generating set. Define the free product of $G_1, \ldots, G_n$ to be the group denoted by $G_1 * \cdots * G_n$ with presentation $\left\langle \bigcup_{i=1}^{n} \mathcal{S}_i \;\middle|\; \bigcup_{i=1}^{n} \mathcal{R}_i \right\rangle$. The concepts of the empty word, syllables, syllable decomposition, and syllable length can be generalized to this arbitrary finite free product of groups setting.

We will now show that if $(G_1, \mathcal{S}_1), (G_2, \mathcal{S}_2), ..., (G_n, \mathcal{S}_n)$ are each totally reflected systems, for $n \geq 2$, then $(G_1 * G_2 * \cdots * G_n, \mathcal{S}_1 \cup \mathcal{S}_2 \cup \cdots \cup \mathcal{S}_n)$ is a totally reflected system. We will first prove this result in the case when $n = 2$. The result for a general finite free product will then follow easily by induction.

Theorem 5.19. *If G and H are totally reflected groups, then the free product $G * H$ is a totally reflected group.*

Proof. Let $\mathcal{S}$ and $\mathcal{T}$ be nice generating sets for G and H, respectively, such that $(G, \mathcal{S})$ and $(H, \mathcal{T})$ are totally reflected systems. We will show that the free product $G * H$ is totally reflected by showing that $(G * H, \mathcal{S} \cup \mathcal{T})$ is a totally reflected system. For convenience, we will abbreviate $\Gamma(G, \mathcal{S})$, $\Gamma(H, \mathcal{T})$, and $\Gamma(G * H, \mathcal{S} \cup \mathcal{T})$ as $\Gamma(G)$, $\Gamma(H)$, and $\Gamma(G * H)$ respectively.

Let's begin by observing that $\Gamma(G)$ and $\Gamma(H)$ embed naturally in $\Gamma(G * H)$ via the standard inclusion functions $\mathfrak{i}_G : \Gamma(G) \to \Gamma(G * H)$ and $\mathfrak{i}_H : \Gamma(H) \to \Gamma(G * H)$. Formally speaking, this means that for any $g \in G$, $h \in H$, $s \in \mathcal{S}^{\pm}$, and $t \in \mathcal{T}^{\pm}$, we have the following:

$$\mathfrak{i}_G(g) = g, \quad \mathfrak{i}_H(h) = h, \quad \mathfrak{i}_G\big(e(g, s)\big) = e(g, s), \quad \mathfrak{i}_H\big(e(h, t)\big) = e(h, t)$$

Note that $\mathfrak{i}_G\big(\Gamma(G)\big)$ and $\mathfrak{i}_H\big(\Gamma(H)\big)$ are subgraphs of $\Gamma(G * H)$. For convenience, denote these subgraphs as Γ_G and Γ_H, respectively. The inclusion functions $\mathfrak{i}_G$ and $\mathfrak{i}_H$ are graph *automorphisms* from $\Gamma(G)$ to Γ_G and from $\Gamma(H)$ to Γ_H, respectively.

Let $s \in \mathcal{S}$ be arbitrary. Consider the edge $e(1, s)$ in $\Gamma(G * H)$. Then $e(1, s)$ is an edge in the subgroup Γ_G, and so $\mathfrak{i}_G^{-1}(e(1, s)) = e(1_G, s)$ is an edge in $\Gamma(G)$. Since $(G, \mathcal{S})$ is a totally reflected system, there is a color-preserving graph reflection on $\Gamma(G)$ of the form $L_s \alpha$, with $\alpha \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$ satisfying $\alpha(s) = s^{-1}$ and $\alpha^2 = \mathrm{id}_G$, which inverts the edge $e(1_G, s)$.

Now, α is a bijective function on G, but we can extend it to a function

α^* on all of $G * H$ by defining $\alpha^*(s) = \alpha(s)$ for all $s \in \mathcal{S}$, by defining $\alpha^*(t) = t$ for all $t \in \mathcal{T}$, and by extending this definition linearly so that it applies to all elements of $G * H$. If $r \in \mathcal{R}_G$, then $\alpha^*(r) = \alpha(r) = 1_G$, since $r = 1_G$ and α is an automorphism on G. If $r \in \mathcal{R}_H$, then $\alpha^*(r) = r = 1_H$. Therefore, α^* respects every relation in $\mathcal{R}_G \cup \mathcal{R}_H$.

We can easily observe that $(\alpha^*)^2 = \mathrm{id}_{G*H}$, since $\alpha^2 = \mathrm{id}_G$. So $(\alpha^*)^2(s) = \alpha^2(s) = s$ for any $s \in \mathcal{S}$ and $(\alpha^*)^2(t) = t$ for any $t \in \mathcal{T}$. Therefore, for any $x \in G * H$, $(\alpha^*)^2(x) = x$, since x can be written as a word in $\mathcal{S}^{\pm} \cup \mathcal{T}^{\pm}$. From this we can conclude that α^* is bijective and thus is an automorphism. Moreover, $(\alpha^*)^{-1} = \alpha^*$. We can easily see that $\alpha^*(\mathcal{S}^{\pm}) = \alpha(\mathcal{S}^{\pm}) = \mathcal{S}^{\pm}$ and $\alpha^*(\mathcal{T}^{\pm}) = \mathcal{T}^{\pm}$. Thus, $\alpha^* \in \mathrm{Aut}(G * H, \mathcal{S}^{\pm} \cup \mathcal{T}^{\pm})$, and by Theorem 3.8 we know that $L_s\alpha^*$ is a color-preserving graph automorphism on $\Gamma(G * H)$. Observe that $L_s\alpha^*(1) = s\alpha^*(1) = s$ and $L_s\alpha^*(s) = s\alpha(s) = ss^{-1} = 1$, implying that $L_s\alpha^*$ inverts the edge $e(1, s)$ in $\Gamma(G * H)$. Our next goal is to show that $L_s\alpha^*$ is a reflection on $\Gamma(G * H)$.

First, we must show that $L_s\alpha^*$ has order 2. To do this, we need only show that $(L_s\alpha^*)^2$ fixes all vertices of $\Gamma(G * H)$. Since $\Gamma(G * H)$ is a simplicial graph, fixing all vertices fixes the entire graph.

Notice that for any $s \in \mathcal{S}$, $(L_s\alpha^*)^2(s) = s\alpha^*(s\alpha^*(s)) = s\alpha^*(s)(\alpha^*)^2(s) = s\alpha(s)1 = ss^{-1} = 1$. Additionally, observe that for any $t \in \mathcal{T}$, $(L_s\alpha^*)^2(t) = s\alpha^*(s\alpha^*(t)) = s\alpha^*(s)(\alpha^*)^2(t) = s\alpha(s)t = ss^{-1}t = 1t = t$. Therefore, $(L_s\alpha^*)^2$ fixes every generator of $G * H$ and thus must fix all elements of $G * H$. Consequently, $(L_s\alpha^*)^2 = \mathrm{id}_{\Gamma(G*H)}$.

Second, we need to identify edges in $\Gamma(G * H)$ which are inverted by $L_s\alpha^*$. To this end, let $\mathcal{E} = \{\text{edges in } \Gamma(G) \text{ which are inverted by } L_s\alpha\} = \Gamma(G)_{L_s\alpha}$

and let $\widetilde{\mathcal{E}} = \mathfrak{i}_G(\mathcal{E})$. Suppose $\widetilde{e} \in \widetilde{\mathcal{E}}$. Then there exists $e \in \mathcal{E}$ such that $\mathfrak{i}_G(e) = \widetilde{e}$. Since e is an edge in $\Gamma(G)$, there exists $g \in G$ and $s \in \mathcal{S}^{\pm}$ with $e = e(g,s)$. So $\widetilde{e} = \mathfrak{i}_G(e(g,s)) = e(g,s)$ also, but as an edge in $\Gamma(G * H)$. Using the fact that $e(g,s)$ is inverted by $L_s\alpha$, we can observe the following:

$$L_s\alpha^* \left(\iota(\widetilde{e})\right) = s\alpha^*(g) = s\alpha(g) = L_s\alpha(g) = gs = \tau(\widetilde{e})$$

$$L_s\alpha^* \left(\tau(\widetilde{e})\right) = L_s\alpha^* \left(L_s\alpha^* \left(\iota(\widetilde{e})\right)\right) = \left(L_s\alpha^*\right)^2 \left(\iota(\widetilde{e})\right) = \iota(\widetilde{e})$$

Therefore, $L_s\alpha^*$ inverts the edge $\widetilde{e} \in \widetilde{\mathcal{E}}$. Since $\widetilde{e}$ was arbitrary, we now know that $L_s\alpha^*$ inverts every edge in $\widetilde{\mathcal{E}}$.

Finally, we must show that the edges of $\widetilde{\mathcal{E}}$, if removed, separate the graph $\Gamma(G * H)$. We observed previously that $L_s\alpha^*$ inverts the edge $e(1,s)$ in $\Gamma(G * H)$. Therefore, in order to show that the set $\widetilde{\mathcal{E}}$ separates $\Gamma(G * H)$, we must show that any path in $\Gamma(G * H)$ from 1 to s contains at least one edge from $\widetilde{\mathcal{E}}$. Suppose, to the contrary, that this is not true. Then the following set must be nonempty:

$$P = \{\text{paths in } \Gamma(G * H) \text{ from 1 to } s \text{ containing no edges from set } \widetilde{\mathcal{E}}\} \neq \varnothing$$

Let γ be an element of P such that $w(\gamma)$, the word in $\mathcal{S}^{\pm} \cup \mathcal{T}^{\pm}$ associated with γ, has shortest syllable length, meaning that $l\big(w(\gamma)\big) \leq l\big(w(\delta)\big)$ where δ is any path in set P.

We can express $w = w(\gamma)$ in terms of its syllable decomposition as follows:

$$w = w_1 w_2 \cdots w_n$$

where each w_i is an element of G or H; $n = l\big(w(\gamma)\big)$; and consecutive syllables w_i and w_{i+1} are not both elements of the same factor group (G or H) in the event that $n \geq 2$. Also, each w_i can be expressed as follows:

$$w_i = \begin{cases} s_1 s_2 \cdots s_{n_i}, & s_k \in \mathcal{S}^{\pm} \text{ for all } k \in I_{n_i} \quad \text{when } w_i \in G \\ t_1 t_2 \cdots t_{n_i}, & t_k \in \mathcal{T}^{\pm} \text{ for all } k \in I_{n_i} \quad \text{when } w_i \in H \end{cases}$$

where n_i is some positive integer. Since γ is a path from 1 to s, we must have that $w = s$. Clearly, s has syllable length of 1 in $G * H$, and therefore so must w. Then we must have that either $n = 1$ or that $n \geq 2$ and $w_{i'} = 1_G$ or 1_H for some $i' \in I_n$.

First consider the case where $n = 1$. Then the syllable decomposition of w consists of a single syllable, w_1. Since γ is a path from 1 to s, where $s \in \mathcal{S}$, w_1 must be an $\mathcal{S}$-syllable. Then w_1 can be expressed as a word in $\mathcal{S}^{\pm}$, $s_1 s_2 \cdots s_{n_1}$, where $s_k \in \mathcal{S}^{\pm}$ for all $k \in I_{n_1}$ and n_1 is some positive integer. The path γ is then completely contained in the Γ_G subgraph of $\Gamma(G * H)$, implying that $\mathfrak{i}_G^{-1}(\gamma) = \widetilde{\gamma}$ is a path in $\Gamma(G)$. Moreover, since γ is a path from 1 to s in $\Gamma(G * H)$, $\widetilde{\gamma}$ is a path from 1_G to s in $\Gamma(G)$. Since $L_s\alpha$ is a reflection on $\Gamma(G)$ which inverts the edge $e(1_G, s)$, any path in $\Gamma(G)$ from 1_G to s must contain at least one edge from the set $\mathcal{E}$. In particular, the path $\widetilde{\gamma}$ in $\Gamma(G)$ must contain at least one edge, say e', from the set $\mathcal{E}$. Then $\mathfrak{i}_G(e')$ is an element of $\widetilde{\mathcal{E}}$. But $\mathfrak{i}_G(e')$ is an edge in the path γ, which gives us a contradiction. We were assuming that γ contained no edges from $\widetilde{\mathcal{E}}$. Thus, this case cannot occur.

Now consider the case where $n \geq 2$. Then for some $i' \in I_n$ we must have that $w_{i'} = 1_G$ or 1_H. Suppose that we remove the syllable $w_{i'}$ from the word

w. Call this newly formed word $\widetilde{w}$, and let $\widetilde{\gamma}$ be the subpath of γ in $\Gamma(G*H)$ with $w(\widetilde{\gamma}) = \widetilde{w}$. As group elements of $G*H$, $w = \widetilde{w}$. Therefore, $w = s$ implies that $\widetilde{w} = s$. So $\widetilde{\gamma}$ is also a path in $\Gamma(G*H)$ from 1 to s. Since γ contains no edges from set $\widetilde{\mathcal{E}}$, we also must have that $\widetilde{\gamma}$ contains no edges from set $\widetilde{\mathcal{E}}$. Therefore, $\widetilde{\gamma}$ is a path in set P. If $i' = 1$ or $i' = n$, then one can easily observe that $l(\widetilde{w}) = n-1$. If $n \geq 3$ and $i' \notin \{1, n\}$, then $w_{i'-1}$ and $w_{i'+1}$ are are both $\mathcal{S}$-syllables or both $\mathcal{T}$-syllables. Then the product $w_{i'-1}w_{i'+1}$ forms a single syllable. Therefore, $l(\widetilde{w}) = n-2$. In any event, we have shown that $l(\widetilde{w}) < n$, where $n = l(w)$. But this is a contradiction! We assumed that γ was an element of P such that $l\big(w(\gamma)\big) = l(w) \leq l\big(w(\delta)\big)$ for any path δ from set P. Thus, this case cannot occur either.

We have now shown that we cannot have $n = 1$ or $n \geq 2$. Thus, such a path γ cannot exist, meaning that set P must, in fact, be empty. Therefore, any path in $\Gamma(G*H)$ from 1 to s must contain at least one edge from the set $\widetilde{\mathcal{E}}$. Consequently, $\widetilde{\mathcal{E}}$ separates $\Gamma(G*H)$.

Therefore, $\widetilde{\mathcal{E}}$ is a set of edges in $\Gamma(G*H)$, each of which is inverted by $L_s\alpha^*$, and $\Gamma(G*H) - \widetilde{\mathcal{E}}$ is disconnected. So $L_s\alpha^*$ is indeed a color-preserving reflection on $\Gamma(G*H)$ which inverts the specified edge $e(1,s)$. A symmetric argument would show that there exists a color-preserving reflection on $\Gamma(G*H)$ inverting any edge of the form $e(1,t)$, where $t \in \mathcal{T}$. Therefore, we have shown that for any $u \in \mathcal{S} \cup \mathcal{T}$, there exists a color-preserving reflection on $\Gamma(G*H)$ inverting the edge $e(1,u)$. Then by Proposition 5.5 we can conclude that $(G*H, \mathcal{S}^{\pm} \cup \mathcal{T}^{\pm})$ is a totally reflected system, implying that $G*H$ is a totally reflected group.

$\square$

Corollary 5.20. *If* $G_1, G_2, ..., G_n$ *are totally reflected groups for* $n \geq 2$*, then the free product* $G_1 * G_2 * \cdots * G_n$ *is a totally reflected group.*

Chapter 6

Right-Angled Products

6.1 Defining a Right-Angled Product

Suppose that $\mathcal{G} = \{G_i \bigm| i \in I_n\}$ is a collection of groups and that Δ is a simplicial graph on n vertices $\{v_i \bigm| i \in I_n\}$. For each $i \in I_n$, let 1_{G_i} denote the identity element of group G_i.

We define a **word in $\mathcal{G}$** to be a sequence, w, of nonnegative-integer length m, where each term in the sequence is an element of G_i for some $i \in I_n$. The **length of a word** w, which we will denote by $|w|$, is defined to be the length of the sequence defining w. The **empty word**, which we will denote by 1, is the unique word of length 0 corresponding to the empty sequence.

We will use juxtaposition notation, rather than comma-separated notation, when working with words. Therefore, any word w of positive length m can be expressed in the form $w = g_1 \cdots g_m$, where each $g_i \in G_j$ for some $j \in I_n$. The terms $g_1, \dots, g_m$ in the word w are called the **syllables** of w. With this new terminology, we can see that $|w|$ equals the number of syllables in w. Consequently, the empty word is said to have zero syllables. Suppose that

$w = g_1 \cdots g_m$ is a nonempty word in $\mathcal{G}$. Then $m \geq 1$. A **subword** of w is any word formed by removing the first i and the last j syllables of w, for some nonnegative integers i and j with $0 \leq i + j \leq m$.

Let $\mathcal{W}$ denote the set of all words in $\mathcal{G}$. We can define a binary operation on $\mathcal{W}$ in terms of concatenation of words. Formally speaking, for any nonempty words $w_1, w_2 \in \mathcal{W}$ with $w_1 = g_1 \cdots g_{m_1}$ and $w_2 = g'_1 \cdots g'_{m_2}$ we define the binary operation, $\bullet$, as follows:

$$w_1 \bullet w_2 := g_1 \cdots g_{m_1} g'_1 \cdots g'_{m_2}$$

Additionally, the concatenation of a word $w \in \mathcal{W}$ and the empty word is defined as follows:

$$1 \bullet w := w \qquad \text{and} \qquad w \bullet 1 := w$$

We often will omit the operation symbol when concatenating words and will write $w_1 w_2$ in place of $w_1 \bullet w_2$. We can easily see that the binary operation of concatenation is associative.

For the sake of clarity, here and throughout this chapter, we will use the symbol $\overset{G_i}{=}$ to indicate that two elements $g, g' \in G_i$ represent the same group element of G_i. Also, we will use the symbol $\overset{\mathcal{W}}{=}$ to indicate that two words w and w' are identical as elements of $\mathcal{W}$. That is, if $w = g_1 \cdots g_m \in \mathcal{W}$ and $w' = h_1 \cdots h_l \in \mathcal{W}$, then

$$w \overset{\mathcal{W}}{=} w' \iff l = m \text{ and for each } i \in I_m \text{ we have } h_i \overset{G_j}{=} g_i \text{ for some } j \in I_m$$

Juxtaposition will be understood to mean concatenation of words when used with the symbol $\overset{\mathcal{W}}{=}$ and will be understood to mean product of group elements under the binary operation of group G_i when used with the symbol $\overset{G_i}{=}$.

In addition to the binary operation on $\mathcal{W}$, we can define some transformations, or **moves**, that can be performed on individual elements of $\mathcal{W}$. We will summarize these moves in the following table. The first four moves are mentioned in [10].

Move	Definition
D_1	remove a syllable g_i such that $g_i \overset{G_j}{=} 1_{G_j}$ for some $j \in I_n$
U_1	insert a syllable of the form 1_{G_j} for some $j \in I_n$
D_2	consolidate adjacent syllables g_i, g_{i+1} which come from the same factor group G_j into a single syllable $h \in G_j$ such that $h \overset{G_j}{=} g_i g_{i+1}$
U_2	split up a syllable g_i coming from G_j into two adjacent syllables h, h' (in that order) with $h, h' \in G_j$ and $g_i \overset{G_j}{=} hh'$
T	transpose adjacent syllables g, g' such that $g \in G_i$, $g' \in G_j$, $i \neq j$, and v_i, v_j are adjacent in Δ

Whenever a word w' can be obtained from a word w using a single move of type M, where $M \in \{D_1, U_1, D_2, U_2, T\}$, we will symbolize this relationship as shown here:

$$w \xrightarrow{M} w'$$

The following relationships can be easily deduced from the definitions of moves D_1, U_1, D_2, U_2, and T:

$$w \xrightarrow{D_1} w' \iff w' \xrightarrow{U_1} w$$

$$w \xrightarrow{D_2} w' \iff w' \xrightarrow{U_2} w$$

$$w \xrightarrow{T} w' \iff w' \xrightarrow{T} w$$

We can define a relation $\sim$ on $\mathcal{W}$ in terms of the moves just discussed. For any $w, w' \in \mathcal{W}$, we'll state that $w \sim w'$ if and only if condition (1) or (2) below is satisfied:

(1) $w \overset{\mathcal{W}}{=} w'$

(2) $w \overset{\mathcal{W}}{\neq} w'$, but there exists $m \in \mathbb{N}$, $m \geq 2$, and a collection of words $\{w_i \mid i \in I_m\}$ such that

 (i) $w_1 \overset{\mathcal{W}}{=} w$

 (ii) $w_m \overset{\mathcal{W}}{=} w'$

 (iii) for each $i \in I_{m-1}$ there exists $M_i \in \{D_1, U_1, D_2, U_2, T\}$ such that $w_i \xrightarrow{M_i} w_{i+1}$

Condition (2) can be summarized by saying that w can be **transformed** into w' using a finite sequence of moves from the set $\{D_1, U_1, D_2, U_2, T\}$.

Lemma 6.1. *The relation $\sim$ is an equivalence relation on $\mathcal{W}$.*

Proof. Suppose that $w \in \mathcal{W}$ is arbitrary. It is clear, by the definition of $\sim$, that $w \sim w$. Therefore, $\sim$ is reflexive.

Now suppose $w, w' \in \mathcal{W}$ are arbitrary. Assume $w \sim w'$. If $w \overset{\mathcal{W}}{=} w'$, then of course $w' \sim w$. Now assume $w \overset{\mathcal{W}}{\neq} w'$. Then $w \sim w'$ implies that there exists $m \in \mathbb{N}$, $m \geq 2$, and a collection of words $\{w_i \mid i \in I_m\}$ such that $w_1 \overset{\mathcal{W}}{=} w$; $w_m \overset{\mathcal{W}}{=} w'$; and for each $i \in I_{m-1}$ there exists $M_i \in \{D_1, U_1, D_2, U_2, T\}$ such that $w_i \xrightarrow{M_i} w_{i+1}$. Symbolically, we can summarize the scenario as shown here:

$$w \overset{\mathcal{W}}{=} w_1 \xrightarrow{M_1} w_2 \xrightarrow{M_2} w_3 \cdots w_{m-1} \xrightarrow{M_{m-1}} w_m \overset{\mathcal{W}}{=} w'$$

For each $i \in I_m$, let $v_i := w_{m-i+1}$. For each $i \in I_{m-1}$, define L_i as follows:

$$L_i := \begin{cases} D_1, & \text{if } M_{m-i} = U_1 \\ U_1, & \text{if } M_{m-i} = D_1 \\ D_2, & \text{if } M_{m-i} = U_2 \\ U_2, & \text{if } M_{m-i} = D_2 \\ T, & \text{if } M_{m-i} = T \end{cases}$$

Then $v_1 \overset{\mathcal{W}}{=} w_m \overset{\mathcal{W}}{=} w'$; $v_m \overset{\mathcal{W}}{=} w_1 \overset{\mathcal{W}}{=} w$; and for each $i \in I_{m-1}$, $v_i \xrightarrow{L_i} v_{i+1}$. Symbolically, we can summarize this scenario as shown here:

$$w' \overset{\mathcal{W}}{=} v_1 \xrightarrow{L_1} v_2 \xrightarrow{L_2} v_3 \cdots v_{m-2} \xrightarrow{L_{m-2}} v_{m-1} \xrightarrow{L_{m-1}} v_m \overset{\mathcal{W}}{=} w$$

We've now shown that $w' \sim w$. Therefore, $\sim$ is symmetric.

Finally, suppose $w, w', w'' \in \mathcal{W}$ are arbitrary. Assume $w \sim w'$ and $w' \sim w''$. If $w \overset{\mathcal{W}}{=} w''$, then of course $w \sim w''$. So assume $w \overset{\mathcal{W}}{\neq} w''$. If $w \overset{\mathcal{W}}{=} w'$, then $w' \sim w''$ implies that $w \sim w''$. If $w' \overset{\mathcal{W}}{=} w''$, then $w \sim w'$ implies that $w \sim w''$. So assume now that $w \overset{\mathcal{W}}{\neq} w'$ and $w' \overset{\mathcal{W}}{\neq} w''$. Then $w \sim w'$ implies that there exists $m \in \mathbb{N}$, $m \geq 2$, and a collection of words $\{w_i \mid i \in I_m\}$ such that $w_1 \overset{\mathcal{W}}{=} w$; $w_m \overset{\mathcal{W}}{=} w'$; and for each $i \in I_{m-1}$ there exists $M_i \in \{D_1, U_1, D_2, U_2, T\}$ such that $w_i \xrightarrow{M_i} w_{i+1}$. Also, $w' \sim w''$ implies that there exists $l \in \mathbb{N}$, $l \geq 2$, and a collection of words $\{v_i \mid i \in I_l\}$ such that $v_1 \overset{\mathcal{W}}{=} w'$; $v_l \overset{\mathcal{W}}{=} w''$; and for each $i \in I_{l-1}$ there exists $L_i \in \{D_1, U_1, D_2, U_2, T\}$ such that $v_i \xrightarrow{L_i} v_{i+1}$. Note that $w_m \overset{\mathcal{W}}{=} w' \overset{\mathcal{W}}{=} v_1$.

For each $i \in I_{l+m}$, let u_i be defined as follows:

$$u_i := \begin{cases} w_i, & \text{if } i \in I_m \\ v_{i-m+1}, & \text{if } i \in I_{l+m-1} - I_{m-1} \end{cases}$$

Since $w_m \overset{\mathcal{W}}{=} v_1$, u_m is well-defined. For each $i \in I_{l+m-2}$, define K_i as follows:

$$K_i := \begin{cases} M_i, & \text{if } i \in I_{m-1} \\ L_{i-m+1}, & \text{if } i \in I_{l+m-2} - I_{m-1} \end{cases}$$

Then $u_1 \overset{\mathcal{W}}{=} w_1 \overset{\mathcal{W}}{=} w$; $u_{l+m-1} \overset{\mathcal{W}}{=} v_l \overset{\mathcal{W}}{=} w''$; and for each $i \in I_{l+m-2}$, $u_i \xrightarrow{K_i} u_{i+1}$. Symbolically, we can summarize the scenario using the diagrams shown below. The first two diagrams show the idea behind the construction of the third diagram.

$$w \overset{\mathcal{W}}{=} w_1 \xrightarrow{M_1} w_2 \cdots w_{m-1} \xrightarrow{M_{m-1}} w_m$$

$$v_1 \xrightarrow{L_1} v_2 \ldots v_{l-1} \xrightarrow{L_{l-1}} v_l \overset{\mathcal{W}}{=} w''$$

$$w \overset{\mathcal{W}}{=} u_1 \xrightarrow{K_1} u_2 \cdots u_{m-1} \xrightarrow{K_{m-1}} u_m \xrightarrow{K_m} u_{m+1} \cdots u_{l+m-2} \xrightarrow{K_{l+m-2}} u_{l+m-1} \overset{\mathcal{W}}{=} w''$$

We've now shown that $w \sim w''$. Therefore, $\sim$ is transitive.

□

For any $w \in \mathcal{W}$, let $[w]$ denote the equivalence class of w with respect to the relation $\sim$. That is, $[w] = \{w' \in \mathcal{W} \mid w \sim w'\}$. Let $\mathcal{W}/\sim$ denote the set of all the equivalence classes under the relation $\sim$. We can define a binary operation, $\blacksquare$, on $\mathcal{W}/\sim$ by stating that $[w]\blacksquare[v] := [w \bullet v] = [wv]$.

Lemma 6.2. *The binary operation* $\blacksquare$ *on* $\mathcal{W}/\sim$ *is well-defined.*

Proof. Suppose that $w, w', v, v' \in \mathcal{W}$. Assume $w \sim w'$ and $v \sim v'$.

We will begin by showing that $wv \sim w'v$. If $w \overset{\mathcal{W}}{=} w'$, then $wv \overset{\mathcal{W}}{=} w'v$. Then by the reflexive property of $\sim$, $wv \sim w'v$. So now assume $w \overset{\mathcal{W}}{\neq} w'$. Then $w \sim w'$ implies that there exists $m \in \mathbb{N}$, $m \geq 2$, and a collection of words $\{w_i \mid i \in I_m\}$ such that $w_1 \overset{\mathcal{W}}{=} w$; $w_m \overset{\mathcal{W}}{=} w'$; and for each $i \in I_{m-1}$, there exists $M_i \in \{D_1, U_1, D_2, U_2, T\}$ such that $w_i \xrightarrow{M_i} w_{i+1}$.

Consider the collection of words $\{w_i v \mid i \in I_m\}$. Note that $w_1 v \overset{\mathcal{W}}{=} wv$, since $w_1 \overset{\mathcal{W}}{=} w$. Also, $w_m v \overset{\mathcal{W}}{=} w'v$, since $w_m \overset{\mathcal{W}}{=} w'$. We can also see that $w_i \xrightarrow{M_i} w_{i+1}$ implies that $w_i v \xrightarrow{M_i} w_{i+1} v$. Since w_i and w_{i+1} are subwords of $w_i v$ and $w_{i+1} v$, respectively, the move that transformed w_i into w_{i+1} can be applied in the same way to the w_i subword of $w_i v$ to obtain $w_{i+1} v$. Therefore, $wv \sim w'v$.

A symmetric argument to the one just performed would allow us to show that $w'v \sim w'v'$. Then by the transitive property of $\sim$, we can conclude that $wv \sim w'v'$. Consequently, if $[w] = [w']$ and $[v] = [v']$, then $[w]\blacksquare[v] = [w']\blacksquare[v']$. Therefore, the binary operation $\blacksquare$ is well-defined and independent of equivalence-class representative. □

Since $\blacksquare$ is well-defined, we can now prove properties of $\blacksquare$ without having to address the issue of equivalence-class-representative independence.

Theorem 6.3. *$\mathcal{W}/\!\sim$ is a group with the binary operation $\blacksquare$.*

Proof. For any $w, v \in \mathcal{W}$, $wv \in \mathcal{W}$. As such, $[w]\blacksquare[v] = [wv]$ is an element of $\mathcal{W}/\!\sim$. Therefore, $\mathcal{W}/\!\sim$ is closed under the operation $\blacksquare$.

The equivalence class of the empty word satisfies the conditions of an identity element. This follows naturally from the way we defined the concatenation of a word with the empty word. Suppose $w \in \mathcal{W}$ is arbitrary. Then observe the following:

$$[1]\blacksquare[w] = [1w] = [w] \text{ and } [w]\blacksquare[1] = [w1] = [w]$$

Therefore, $[1]$ will serve as the identity element of $\mathcal{W}/\!\sim$ under the binary operation $\blacksquare$.

The associativity of $\blacksquare$ follows easily from the associativity of the operation of word concatenation. To see this, suppose that $w, v, u \in \mathcal{W}$. Then observe the following:

$$\Big([w]\blacksquare[v]\Big)\blacksquare[u] = [wv]\blacksquare[u] = [(wv)u] = [w(vu)] = [w]\blacksquare[vu] = [w]\blacksquare\Big([v]\blacksquare[u]\Big)$$

Therefore, the binary operation $\blacksquare$ on $\mathcal{W}/\!\sim$ is associative.

Finally, we will show that for any $[w] \in \mathcal{W}/\!\sim$, there exists $[v] \in \mathcal{W}/\!\sim$ such that $[w]\blacksquare[v] = [1]$ and $[v]\blacksquare[w] = [1]$. To this end, suppose $w \in \mathcal{W}$ is arbitrary. If $w \overset{\mathcal{W}}{=} 1$, then $[w]$ can serve as its own inverse element, since $[w]\blacksquare[w] = [1]\blacksquare[1] = [11] = [1]$ So assume $w \overset{\mathcal{W}}{\neq} 1$. Suppose $w = g_1 \cdots g_m$. For each $i \in I_m$, let $G_{f(i)}$ denote the factor group containing g_i. Let $v = g_m^{-1} \cdots g_1^{-1}$, where g_i^{-1} is understood to be the inverse element of g_i in $G_{f(i)}$ for each $i \in I_m$. We claim that $wv \sim 1$. The following diagram illustrates the proof of this claim.

$$\begin{aligned}
g_1 \cdots g_{m-1} g_m g_m^{-1} g_{m-1}^{-1} \cdots g_1^{-1} &\xrightarrow{D_2} g_1 \cdots g_{m-1} 1_{G_{f(m)}} g_{m-1}^{-1} \cdots g_1^{-1} \\
&\xrightarrow{D_1} g_1 \cdots g_{m-1} g_{m-1}^{-1} \cdots g_1^{-1} \\
&\xrightarrow{D_2} g_1 \cdots g_{m-2} 1_{G_{f(m-1)}} g_{m-2}^{-1} \cdots g_1^{-1} \\
&\xrightarrow{D_1} g_1 \cdots g_{m-2} g_{m-2}^{-1} \cdots g_1^{-1} \\
&\vdots \\
&\xrightarrow{D_1} g_1 g_1^{-1} \\
&\xrightarrow{D_2} 1_{G_{f(1)}} \\
&\xrightarrow{D_1} 1
\end{aligned}$$

One can see then that wv can be transformed into the empty word by repeated applications of a D_2 move followed by a D_1 move. Therefore, $wv \sim 1$. A symmetric argument would show that $vw \sim 1$ also. So, $[w] \blacksquare [v] = [wv] = [1]$ and $[v] \blacksquare [w] = [vw] = [1]$. Then for any $[w] \in \mathcal{W}/\sim$, we can define $[w]^{-1}$ as follows:

$$[w]^{-1} := \begin{cases} [w], & \text{if } w \overset{\mathcal{W}}{=} 1 \\ [g_m^{-1} \cdots g_1^{-1}], & \text{if } w \overset{\mathcal{W}}{\neq} 1, w = g_1 \cdots g_m \end{cases}$$

Thus we have proven the existence of inverses in $\mathcal{W}/\sim$ with respect to the operation $\blacksquare$.

Therefore, $\mathcal{W}/\sim$ is a group with binary operation $\blacksquare$.

□

Henceforth, we often will omit the operation symbol ▪ and write $[w][v]$ in place of $[w]\blacksquare[v]$.

We will call the group $\mathcal{W}/\sim$ the **right-angled product determined by $\boldsymbol{\Delta}$ and $\boldsymbol{\mathcal{G}}$**. To emphasize the dependence on the graph Δ and the collection of groups $\mathcal{G}$, we will henceforth denote this group using the symbol $G(\Delta, \mathcal{G})$. Throughout this chapter, the symbols Δ and $\mathcal{G}$ will be understood in the sense established at the beginning of this section. Any exceptions to this convention will be explicitly noted.

6.2 Universal Mapping Property Determined by $\mathcal{G}$ and Δ

Definition 6.4. *Let $\mathcal{G}$ and Δ be as in the preceding section. Suppose G is some fixed group and that $\{\phi_i : G_i \to G \mid i \in I_n\}$ is a collection of homomorphisms satisfying the condition below, which we henceforth will call the **$\boldsymbol{\Delta}$-condition**:*

$$v_i, v_j \text{ adjacent in } \Delta \; (i, j \in I_n, i \neq j)$$
$$\Downarrow$$
$$\phi_i(x)\phi_j(y) = \phi_j(y)\phi_i(x), \; \forall x \in G_i, \forall y \in G_j$$

*We say that G satisfies the **universal mapping property** (or **U.M.P.**) **determined by $\boldsymbol{\mathcal{G}}$ and $\boldsymbol{\Delta}$** with respect to the homomorphisms $\phi_1, \ldots, \phi_n$ if and only if the following condition holds:*

For any group H and any collection of homomorphisms $\{\theta_i : G_i \to H \mid$

$i \in I_n\}$ satisfying the Δ-condition, there exists a unique homomorphism $\psi : G \to H$ such that for any $i \in I_n$, $\theta_i = \psi\phi_i$, that is, the following diagram commutes:

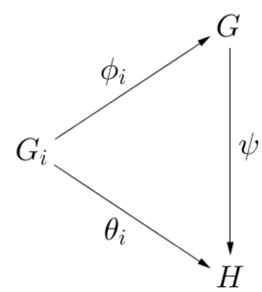

Figure 6.1: Commutative diagram from definition of U.M.P.

We will now show that if a group satisfying the universal mapping property determined by $\mathcal{G}$ and Δ does exist, then that group is unique up to isomorphism.

Proposition 6.5. *Let G and G' be two groups satisfying the universal mapping property determined by $\mathcal{G}$ and Δ, G with respect to $\{\phi_i \mid i \in I_n\}$ and G' with respect to $\{\phi_i' \mid i \in I_n\}$. Then there exists a unique isomorphism $\psi : G \to G'$ such that the following diagram commutes for any $i \in I_n$:*

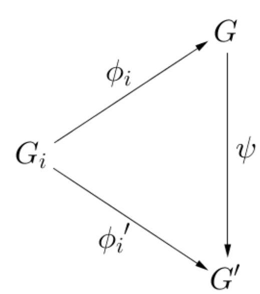

Figure 6.2: Commutative diagram from U.M.P. isomorphism result

Proof. Built in to the fundamental assumptions is the additional supposition coming from Definition 6.4 that both collections of homomorphisms $\{\phi_i \mid i \in I_n\}$ and $\{\phi'_i \mid i \in I_n\}$ satisfy the Δ-condition. Furthermore, by Definition 6.4 there exist unique homomorphisms ψ and ψ' such that the diagrams in Figures 6.3 and 6.4 commute for any $i \in I_n$.

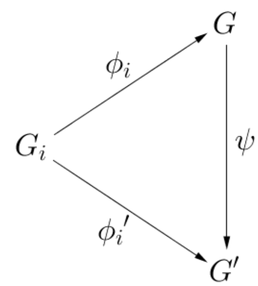

Figure 6.3: Commutative diagram resulting from G satisfying U.M.P. with $H = G'$

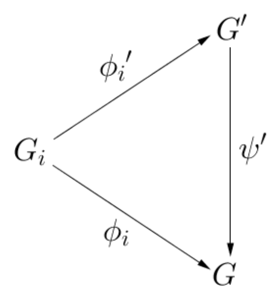

Figure 6.4: Commutative diagram resulting from G' satisfying U.M.P. with $H = G$

Consequently, the compositions $\psi'\psi : G \to G$ and $\psi\psi' : G' \to G'$ are homomorphisms. Also, since the diagrams in Figures 6.3 and 6.4 commute, we know that $\phi_i = \psi'\phi'_i$ and $\phi'_i = \psi\phi_i$. These equalities, together with the associativity of function composition, imply the following:

$$\phi_i = \psi'\phi'_i = \psi'(\psi\phi_i) = (\psi'\psi)\phi_i$$

$$\phi_i' = \psi\phi_i = \psi(\psi'\phi_i') = (\psi\psi')\phi_i'$$

Equivalently, $\psi'\psi : G \to G$ and $\psi\psi' : G' \to G'$ are homomorphisms which, respectively, make the diagrams in Figures 6.5 and 6.6 commute.

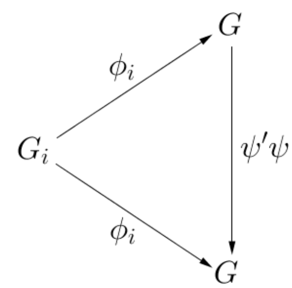

Figure 6.5: Commutative diagram resulting from G satisfying U.M.P. with $H = G$

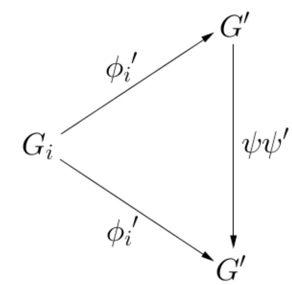

Figure 6.6: Commutative diagram resulting from G' satisfying U.M.P. with $H = G'$

Definition 6.4 tells us that $\psi'\psi$ and $\psi\psi'$ must be the *unique* homomorphisms which, respectively, make the diagrams in Figures 6.5 and 6.6 commute. However, if we replace $\psi'\psi$ and $\psi\psi'$ by id_G and $\mathrm{id}_{G'}$, respectively, the diagrams still commute.

Then by the uniqueness guaranteed by Definition 6.4, we must have that $\psi'\psi = \mathrm{id}_G$ and $\psi\psi' = \mathrm{id}_{G'}$. Therefore, by basic group theory we can conclude that ψ and ψ' are, in fact, inverse isomorphisms. So, $\psi : G \to G'$ is the unique isomorphism which makes the diagram in Figure 6.2 commute.

□

We must now prove the existence of a group satisfying the universal mapping property determined by $\mathcal{G}$ and Δ. Specifically, we will show that the right-angled product $G(\Delta, \mathcal{G})$ is such a group.

Proposition 6.6. *The right-angled product $G(\Delta, \mathcal{G})$ satisfies the universal mapping property determined by $\mathcal{G}$ and Δ.*

Proof. For each $i \in I_n$, define a function $\phi_i : G_i \to G(\Delta, \mathcal{G})$ by stating that $\phi_i(g) = [g]$ for any $g \in G_i$. This function is well-defined since any element of G_i is itself a word in $\mathcal{G}$. The function ϕ_i is, in fact, a homomorphism. To prove this, we must show that ϕ_i respects the group operations of G_i and $G(\Delta, \mathcal{G})$. For the sake of clarity, we will use the binary-operation symbols defined in the preceding section. Additionally, for any group G other than $G(\Delta, \mathcal{G})$, we will use the symbol $\overset{G}{\blacktriangle}$ to denote the binary operation of group G.

Let $g, h \in G_i$. In order to show that ϕ_i respects the group operations of G_i and $G(\Delta, \mathcal{G})$, we must show that $[g \overset{G_i}{\blacktriangle} h] = [g] \blacksquare [h]$. As words in $\mathcal{G}$, g, h, and $g \overset{G_i}{\blacktriangle} h$ are all single-syllable words. By performing a single move of type U_2, $g \overset{G_1}{\blacktriangle} h$ can be transformed into the two-syllable word gh. Symbolically, $g \overset{G_i}{\blacktriangle} h \overset{U_2}{\longrightarrow} gh$. Therefore, $g \overset{G_1}{\blacktriangle} h \sim gh$. Consequently,

$$\phi_i(g \overset{G_i}{\blacktriangle} h) = [g \overset{G_i}{\blacktriangle} h] = [gh] = [g] \blacksquare [h] = \phi_i(g) \blacksquare \phi_i(h)$$

Therefore, for each $i \in I_n$, $\phi_i : G_i \to G(\Delta, \mathcal{G})$, defined as above, is a well-defined homomorphism.

The collection of homomorphisms $\{\phi_i \mid i \in I_n\}$ satisfies the Δ-condition. To see this, suppose v_i, v_j are adjacent in Δ for some $i, j \in I_n$, $i \neq j$. Let $x \in G_i$, $y \in G_j$ be arbitrary. The definition of move T then allows us to observe that $xy \xrightarrow{U_2} yx$. Therefore, $xy \sim yx$ and so the following is true:

$$\phi_i(x)\blacksquare\phi_j(y) = [x]\blacksquare[y] = [xy] = [yx] = [y]\blacksquare[x] = \phi_j(y)\blacksquare\phi_i(x)$$

This shows that the collection of homomorphisms $\{\phi_i \mid i \in I_n\}$ does, indeed, satisfy the Δ-condition.

Suppose H is any group and $\{\theta_i \mid i \in I_n\}$ is any collection of homomorphisms satisfying the Δ-condition. Define a function $\psi : G(\Delta, \mathcal{G}) \to H$ as shown here:

$$\psi([w]) := \begin{cases} 1_H, & \text{if } w \overset{\mathcal{W}}{=} 1 \\ \theta_{f(1)}(g_1) \cdots \theta_{f(m)}(g_m), & \text{if } w \overset{\mathcal{W}}{\neq} 1 \text{ and } w = g_1 \cdots g_m \end{cases}$$

where $f(i)$ indicates the index on the factor group containing g_i, that is, $g_i \in G_{f(i)}$ for each $i \in I_n$.

We need to show that ψ is well-defined and independent of equivalence-class representative. That is, if $[w] = [w']$, we must show that $\psi([w]) = \psi([w'])$. First, we will show that for any $M \in \{D_1, U_1, D_2, U_2, T\}$, $w \xrightarrow{M} w'$ implies that $\psi([w]) = \psi([w'])$.

If w' is obtained from w by performing just a single move, then w and w' differ only by a subword consisting of just one or two syllables. The remaining syllables are unaffected. Since ψ is defined in terms of the action of elements of $\{\theta_i \mid i \in I_n\}$ on individual syllables, we need only prove that the effected syllables behave nicely under the influence of ψ.

Case i: $1_{G_i} \xrightarrow{D_1} 1$ (equivalently, $1 \xrightarrow{U_1} 1_{G_i}$)

$$\psi([1_{G_i}]) = \theta_{f(i)}(1_{G_i}) = 1_H = \psi([1])$$

Case ii: $g_{i-1}1_{G_i}g_{i+1} \xrightarrow{D_1} g_{i-1}g_{i+1}$ (equivalently, $g_{i-1}g_{i+1} \xrightarrow{U_1} g_{i-1}1_{G_i}g_{i+1}$)

$$\begin{aligned}
\psi([g_{i-1}1_{G_i}g_{i+1}]) &= \theta_{f(i-1)}(g_{i-1}) \overset{H}{\blacktriangle} \theta_{f(i)}(1_{G_i}) \overset{H}{\blacktriangle} \theta_{f(i+1)}(g_{i+1}) \\
&= \theta_{f(i-1)}(g_{i-1}) \overset{H}{\blacktriangle} 1_H \overset{H}{\blacktriangle} \theta_{f(i+1)}(g_{i+1}) \\
&= \theta_{f(i-1)}(g_{i-1}) \overset{H}{\blacktriangle} \theta_{f(i+1)}(g_{i+1}) \\
&= \psi([g_{i-1}g_{i+1}])
\end{aligned}$$

Case iii: $g_i g_{i+1} \xrightarrow{D_2} h$, where $g_i, g_{i+1}, h \in G_j$ for some $j \in I_n$ and $g_i \overset{G_j}{\blacktriangle} g_{i+1} \overset{G_j}{=} h$ (equivalently, $h \xrightarrow{U_2} g_i g_{i+1}$)

$$\psi([g_i g_{i+1}]) = \theta_j(g_i) \overset{H}{\blacktriangle} \theta_j(g_{i+1}) = \theta_j(g_i \overset{G_j}{\blacktriangle} g_{i+1}) = \theta_j(h) = \psi([h])$$

Case iv: $gh \xrightarrow{T} hg$, where $j \in G_i$, $h \in G_j$, $i \neq j$, and v_i, v_j adjacent in Δ

$$\psi([gh]) = \theta_i(g)\theta_j(h) = \theta_j(h)\theta_i(g) = \psi([hg])$$

As a result of Cases i-iv, we can conclude that for any move M, $w \xrightarrow{M} w'$ implies that $\psi([w]) = \psi([w'])$.

If, more generally speaking, $w \sim w'$ and $w \overset{\mathcal{W}}{\neq} w'$, then there exists $m \in \mathbb{N}$, $m \geq 2$, and a collection of moves $\{M_i \mid i \in I_m\}$ such that $w_1 \overset{\mathcal{W}}{=} w$; $w_m \overset{\mathcal{W}}{=} w'$; and for each $i \in I_{m-1}$, $w_i \xrightarrow{M_i} w_{i+1}$. Then the above argument allows us to observe the following:

$$\psi([w_1]) = \psi([w_2]) = \cdots = \psi([w_m])$$

Since $w_1 \overset{\mathcal{W}}{=} w$ and $w_m \overset{\mathcal{W}}{=} w$, we can conclude that $\psi([w]) = \psi([w'])$. Therefore, if $[w] = [w']$, then $\psi([w]) = \psi([w'])$.

Let $i \in I_n$ be arbitrary and let $g \in G_i$. Then g is a single-syllable word in $\mathcal{G}$. Therefore, $\psi\phi_i(g) = \psi([g]) = \theta_i(g)$, which proves that the diagram in Figure 6.7 commutes.

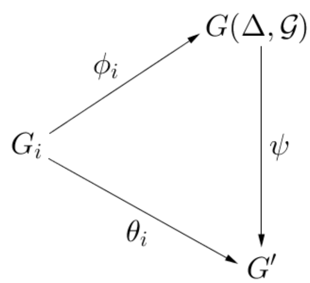

Figure 6.7: Commutative diagram used to show $G(\Delta, \mathcal{G})$ satisfies U.M.P.

Suppose $\psi' : G(\Delta, \mathcal{G}) \to H$ is another homomorphism which makes the diagram in Figure 6.7 commute for any $i \in I_n$. Then for any $i \in I_n$, $\theta_i = \psi'\phi_i$. Let $[w] \in G(\Delta, \mathcal{G}) - \{[1]\}$ be arbitrary. Then $w = g_1 \cdots g_m$. In order for ψ' to be a homomorphism, ψ' must satisfy the following:

$$\begin{aligned}
\psi'([g_1 \cdots g_m]) &= \psi'([g_1] \blacksquare \cdots \blacksquare [g_m]) \\
&= \psi'([g_1]) \overset{H}{\blacktriangle} \cdots \overset{H}{\blacktriangle} \psi'([g_m]) \\
&= \psi'(\phi_{f(1)}(g_1)) \overset{H}{\blacktriangle} \cdots \overset{H}{\blacktriangle} \psi'(\phi_{f(m)}(g_m)) \\
&= \theta_{f(1)}(g_1) \overset{H}{\blacktriangle} \cdots \overset{H}{\blacktriangle} \theta_{f(m)}(g_m)
\end{aligned}$$

$$= \psi\phi_{f(1)}(g_1) \overset{H}{\blacktriangle} \cdots \overset{H}{\blacktriangle} \psi\phi_{f(m)}(g_m)$$

$$= \psi([g_1]) \overset{H}{\blacktriangle} \cdots \overset{H}{\blacktriangle} \psi([g_m])$$

$$= \psi([g_1] \blacksquare \cdots \blacksquare [g_m])$$

$$= \psi([g_1 \cdots g_m])$$

Therefore, for any $[w] \in G(\Delta, \mathcal{G})$ $[w] \overset{\mathcal{W}}{\neq} 1$, $\psi'([w]) = \psi([w])$. And obviously $\psi'([1]) = 1_H = \psi([1])$, since ψ and ψ' are homomorphisms. Therefore, $\psi = \psi'$. So ψ is the *unique* homomorphism which makes the diagram in Figure 6.7 commute.

Therefore, we have shown that $G(\Delta, \mathcal{G})$ satisfies the universal mapping property determined by $\mathcal{G}$ and Δ.

□

6.3 Group Presentation for the Right-Angled Product

The goal of this section is to give an alternate characterization of the right-angled product. Let $\mathcal{G}$ and Δ be as defined earlier in this chapter. If there exists an edge in Δ between two vertices v_i and v_j, we will denote this edge as $\{v_i, v_j\}$.

For each $i \in I_n$, suppose the group G_i is defined in terms of a presentation $\langle \mathcal{S}_i \mid \mathcal{R}_i \rangle$. For any $i, j \in I_n$, $i \neq j$, define a set $\mathcal{R}_{ij}$ of words in $\bigcup_{i \in I_n} \mathcal{S}_i^{\pm}$ as follows:

$$\mathcal{R}_{ij} \coloneqq \{[s,t] \mid s \in \mathcal{S}_i, t \in \mathcal{S}_j\}$$

where $[s,t]$ denotes the commutator $sts^{-1}t^{-1}$. Define a group $\overline{G}(\Delta,\mathcal{G})$ as shown here:

$$\overline{G}(\Delta,\mathcal{G}) \coloneqq \left\langle \bigcup_{i\in I_n} \mathcal{S}_i \;\middle|\; \bigcup_{i\in I_n} \mathcal{R}_i, \bigcup_{\substack{i,j\in I_n \\ \{i,j\}\in E(\Delta)}} \mathcal{R}_{ij} \right\rangle$$

Proposition 6.7. *The group $\overline{G}(\Delta,\mathcal{G})$ is isomorphic to the right-angled product $G(\Delta,\mathcal{G})$.*

Proof. We will show that $\overline{G}(\Delta,\mathcal{G})$ satisfies the universal mapping property determined by $\mathcal{G}$ and Δ. Then Proposition 6.5 will give us the result we seek. For convenience, we will denote $\overline{G}(\Delta,\mathcal{G})$ simply as $\overline{G}$ throughout this proof.

Let $i \in I_n$ be arbitrary. Let $\rho_i : \mathcal{S}_i^{\pm} \to \overline{G}$ be the inclusion function, whereby $\rho_i(s) = s$ for all $s \in \mathcal{S}_i^{\pm}$. Suppose that $r \in \mathcal{R}_i$. Then r is a word in $\mathcal{S}_i^{\pm}$, say $r = s_1 \cdots s_m$. Since r is a relation for the group G_i, $s_1 \cdots s_m \overset{G_i}{=} 1_{G_i}$. Observe the following:

$$\rho_i(s_1)\cdots\rho_i(s_m) = s_1 \cdots s_m \overset{\overline{G}}{=} 1_{\overline{G}}$$

Then from basic theory concerning group presentations, we know that ρ_i can be extended uniquely to a homomorphism $\phi_i : G_i \to \overline{G}$. Since $\phi_i\big|_{\mathcal{S}_i^{\pm}} = \rho_i$ is the inclusion function on $\mathcal{S}_i^{\pm}$ and since $\mathcal{S}_i$ generates G_i, ϕ_i must be the inclusion homomorphism on all of G_i. Since $i \in I_n$ was arbitrary, $\phi_i : G_i \to \overline{G}$ is the inclusion homomorphism for each $i \in I_n$.

Suppose $i,j \in I_n$, $i \neq j$. Let $g \in G_i$, $h \in G_j$ be arbitrary. Assume $\{i,j\} \in E(\Delta)$. The collection of relators $\mathcal{R}_{ij}$ then tells us that each of the elements of $\mathcal{S}_i^{\pm}$ commutes with each of the elements of $\mathcal{S}_j^{\pm}$. Then we also must have that each of the elements of G_i commutes with each of the elements of G_j, since $\mathcal{S}_i$ and $\mathcal{S}_j$ generate G_i and G_j, respectively. Therefore, $gh \overset{\overline{G}}{=} hg$. Observe the following:

$$\phi_i(g)\phi_j(h) = gh = hg = \phi_j(h)\phi_i(g)$$

Thus, the collection of homomorphisms $\{\phi_i : G_i \to \overline{G} \mid i \in I_n\}$ satisfies the Δ-condition.

Suppose H is some fixed group and $\{\theta_i : G_i \to H \mid i \in I_n\}$ is any collection of homomorphisms satisfying the Δ-condition. Let us define a function $\sigma : \bigcup_{i \in I_n} \mathcal{S}_i^{\pm} \to H$. If $s \in \bigcup_{i \in I_n} \mathcal{S}_i^{\pm}$, then $s \in \mathcal{S}_j^{\pm}$ for some $j \in I_n$. Then define σ by stating that $\sigma(s) = \theta_j(s)$. Suppose $r \in \Big(\bigcup_{i \in I_n} \mathcal{R}_i\Big) \cup \Big(\bigcup_{\substack{i,j \in I_n \\ \{i,j\} \in E(\Delta)}} \mathcal{R}_{ij}\Big)$ is arbitrary. Then r is a word in $\bigcup_{i \in I_n} \mathcal{S}_i^{\pm}$, say $r = s_1 \cdots s_m$. If we can show that $\sigma(s_1) \cdots \sigma(s_m) \overset{H}{=} 1_H$, then we can conclude that σ extends to a homomorphism $\psi : \overline{G} \to H$.

Case i: Suppose $r \in \mathcal{R}_j$ for some $j \in I_n$. Then $s_i \in \mathcal{S}_j^{\pm}$ for all $i \in I_m$. Since r is a relation on G_j, $s_1 \cdots s_m \overset{G_j}{=} 1_{G_j}$. Therefore,

$$\sigma(s_1) \cdots \sigma(s_m) = \theta_j(s_1) \cdots \theta_j(s_m) = \theta_j(s_1 \cdots s_m) = \theta_j(1_{G_j}) = 1_H$$

Case ii: Suppose $r \in \mathcal{R}_{ij}$ for some $i, j \in I_n$ with $\{i,j\} \in E(\Delta)$. Then $r = sts^{-1}t^{-1}$ with $s \in \mathcal{S}_i$ and $t \in \mathcal{S}_j$. Therefore,

$$\begin{aligned}
\sigma(s)\sigma(t)\sigma(s^{-1})\sigma(t^{-1}) &= \theta_i(s)\theta_j(t)\theta_i(s^{-1})\theta_j(t^{-1}) \\
&= \theta_i(s)\theta_i(s^{-1})\theta_j(t)\theta_j(t^{-1}) \quad \text{(by the Δ-condition)} \\
&= \theta_i(s)\left(\theta_i(s)\right)^{-1}\theta_j(t)\left(\theta_j(t)\right)^{-1} \\
&= 1_H 1_H = 1_H
\end{aligned}$$

Cases i and ii show that no matter what type of relation r is, $\sigma(s_1) \cdots$

$\sigma(s_m) \overset{H}{=} 1_H$. Therefore, σ can be extended to a homomorphism $\psi : \overline{G} \to H$.

Let $i \in I_n$ be arbitrary. Suppose $g \in G_i$. Since $\mathcal{S}_i$ generates G_i, g can be written as a word in $\mathcal{S}_i^{\pm}$, say $g = s_1 \cdots s_m$. Then observe the following:

$$\psi\left(\phi_i(g)\right) = \psi(g) = \psi(s_1 \cdots s_m) = \psi(s_1) \cdots \psi(s_m) = \sigma(s_1) \cdots \sigma(s_m)$$
$$= \theta_i(s_1) \cdots \theta_i(s_m) = \theta_i(s_1 \cdots s_m) = \theta_i(g)$$

Therefore, $\theta_i = \psi\phi_i$ for each $i \in I_n$. Equivalently, ψ makes the diagram in Figure 6.8 commute for each $i \in I_n$.

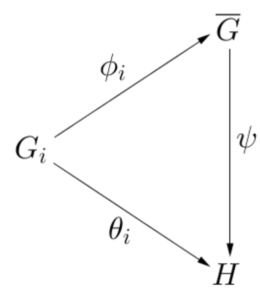

Figure 6.8: Commutative diagram used to show $\overline{G}(\Delta, \mathcal{G})$ satisfies U.M.P.

Suppose that $\psi' : \overline{G} \to H$ is another homomorphism which makes the diagram in Figure 6.8 commute for each $i \in I_n$. Let $g \in \overline{G}$. Then g can be written as a word in $\bigcup_{i \in I_n} \mathcal{S}_i^{\pm}$, say $g = s_1 \cdots s_m$. For each $i \in I_m$, let $f(i)$ denote the index on the factor group containing s_i. That is, $s_i \in G_{f(i)}$, or more specifically, $s_i \in \mathcal{S}_{f(i)}^{\pm}$. Then observe the following:

$$\begin{aligned}\psi'(g) &= \psi'(s_1 \cdots s_m)\\ &= \psi'(s_1) \cdots \psi'(s_m)\\ &= \psi'\left(\phi_{f(1)}(s_1)\right) \cdots \psi'\left(\phi_{f(m)}(s_m)\right)\end{aligned}$$

$$
\begin{aligned}
&= \theta_{f(1)}(s_1)\cdots\theta_{f(m)}(s_m) \\
&= \psi\left(\phi_{f(1)}(s_1)\right)\cdots\psi\left(\theta_{f(m)}(s_m)\right) \\
&= \psi(s_1)\cdots\psi(s_m) \\
&= \psi(s_1\cdots s_m) \\
&= \psi(g)
\end{aligned}
$$

We have shown then that $\psi'(g) = \psi(g)$ for all $g \in \overline{G}$. Therefore, $\psi' = \psi$, implying that ψ is, in fact, the unique homomorphism which makes the diagram in Figure 6.8 commute for each $i \in I_n$.

Therefore, $\overline{G} = \overline{G}(\Delta, \mathcal{G})$ satisfies the universal mapping property determined by $\mathcal{G}$ and Δ, and by Proposition 6.5 we must have that $\overline{G}(\Delta, \mathcal{G}) \cong G(\Delta, \mathcal{G})$.

□

Our original definition of the right-angled product given in Section 6.1 required a great deal of preliminary setup. In the end, however, we were left with a very thorough description of the group $G(\Delta, \mathcal{G})$. The group-presentation approach that we used to create $\overline{G}(\Delta, \mathcal{G})$ provided us with a more compact, minimalistic, easy-to-state definition, but it does have its drawbacks. Any one group may admit several different presentations. There is also the well-known difficulty in determining whether or not groups given by different presentations are isomorphic. These problems aside, the presentation form of the right-angled product, $\overline{G}(\Delta, \mathcal{G})$, gives us a convenient way of stating examples of this newly defined group construction.

6.4 Notable Examples of Right-Angled Products

The right-angled product construction provides a nice generalization of some other well-known group constructions, such as the direct product and free product. Indeed, we will show that the direct product $G_1 \times \cdots \times G_n$ and the free product $G_1 * \cdots * G_n$ are examples of right-angled products, where the graph Δ is chosen in a very particular way in each case. We will also show that right-angled Artin groups can be viewed more generally as right-angled products.

6.4.1 Finite Direct Products

Let $\mathcal{G} = \{G_i \bigm| i \in I_n\}$ and let Δ be the complete graph on n vertices $\{v_i \bigm| i \in I_n\}$. That is, for any $i, j \in I_n$, $i \neq j$, $\{v_i, v_j\} \in E(\Delta)$. Consider the direct product $G_1 \times \cdots \times G_n$. This is the group $\{(g_1, \ldots, g_n) \bigm| g_i \in G_i \ \forall i \in I_n\}$ with binary operation as shown here:

$$(g_1, \ldots, g_n)(h_1, \ldots, h_n) = (g_1 h_1, \ldots, g_n h_n)$$

The identity element of $G_1 \times \cdots \times G_n$ is $(1_{G_1}, \ldots, 1_{G_n})$. We will show that $G_1 \times \cdots \times G_n$ satisfies the universal mapping property determined by $\mathcal{G}$ and Δ.

For any $i \in I_n$, let $\phi_i : G_i \to G_1 \times \cdots \times G_n$ be the standard inclusion function for the direct product. That is, for any $g \in G_i$, $\phi_i(g)$ is the element of $G_1 \times \cdots \times G_n$ created by replacing the i^{th} coordinate of the identity element by g, as shown here:

$$\phi_i(g) := (1_{G_1}, \ldots, 1_{G_{i-1}}, g, 1_{G_{i+1}}, \ldots, 1_{G_n})$$

One can easily observe that this function is a homomorphism. Also, the collection of homomorphisms $\{\phi_i : G_i \to G_1 \times \cdots \times G_n \mid i \in I_n\}$ satisfies the Δ-condition. To see this, suppose that $i, j \in I_n$, $i \neq j$. Then $\{v_i, v_j\} \in E(\Delta)$. Suppose $g \in G_i$, $h \in G_j$ are arbitrary. Then the products $\phi_i(g)\phi_j(h)$ and $\phi_j(h)\phi_i(g)$ both give us the element of $G_1 \times \cdots \times G_n$ created by replacing the i^{th} and j^{th} coordinates of the identity element by g and h, respectively. Therefore, $\phi_i(g)\phi_j(h) = \phi_j(h)\phi_i(g)$, showing that the Δ-condition is satisfied.

Suppose H is some fixed group and $\{\theta_i : G_i \to H \mid i \in I_n\}$ is any collection of homomorphisms satisfying the Δ-condition. That is, for any $i, j \in I_n$, $i \neq j$, and for any $g \in G_i$, $h \in G_j$, $\theta_i(g)\theta_j(h) = \theta_j(h)\theta_i(g)$.

Define a function $\psi : G_1 \times \cdots \times G_n \to H$ as shown here:

$$\psi(g_1, \ldots, g_n) := \theta_1(g_1) \cdots \theta_n(g_n)$$

Suppose $(g_1, \ldots, g_n), (h_1, \ldots, h_n) \in G_1 \times \cdots \times G_n$ are arbitrary. Then observe the following:

$$\begin{aligned}
\psi\left((g_1, \ldots, g_n)(h_1, \ldots, h_n)\right) &= \psi(g_1h_1, \ldots, g_nh_n) \\
&= \theta_1(g_1h_1) \cdots \theta_n(g_nh_n) \\
&= \theta_1(g_1)\theta_1(h_1) \cdots \theta_n(g_n)\theta_n(h_n) \\
&= \theta_1(g_1) \cdots \theta_n(g_n)\theta_1(h_1) \cdots \theta_n(h_n) \\
&= \psi(g_1, \ldots, g_n)\psi(h_1, \ldots, h_n)
\end{aligned}$$

Therefore, ψ is a homomorphism.

Fix $i \in I_n$ and let $g \in G_i$ be arbitrary. Then $\phi_i(g)$ is the n-tuple created by replacing the i^{th} coordinate of the identity element with g. Since $\theta_j(1_{G_j}) = 1_H$ for all $j \in I_n$, we must then have that $\psi(\phi_i(g)) = \theta_i(g)$. Therefore, ψ makes the diagram in Figure 6.9 commute for each $i \in I_n$, and one can easily check that ψ is the *unique* homomorphism with this property. So, $G_1 \times \cdots \times G_n$ satisfies the universal mapping property determined by $\mathcal{G}$ and Δ, and by Proposition 6.5, we must have that $G_1 \times \cdots \times G_n \cong G(\Delta, \mathcal{G})$, where Δ is the complete graph on n vertices.

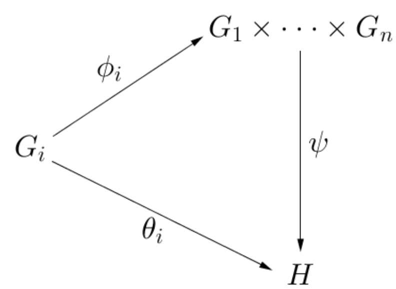

Figure 6.9: Commutative diagram used to show $G_1 \times \cdots \times G_n$ satisfies U.M.P.

If each group G_i is given by a presentation $\langle \mathcal{S}_i \mid \mathcal{R}_i \rangle$, then it is well-known that the direct product $G_1 \times \cdots \times G_n$ is isomorphic to the group given by the following presentation:

$$\left\langle \bigcup_{i \in I_n} \mathcal{S}_i \;\middle|\; \bigcup_{i \in I_n} \mathcal{R}_i, \bigcup_{\substack{i,j \in I_n \\ i \neq j}} \mathcal{R}_{ij} \right\rangle$$

where, recall, $\mathcal{R}_{ij}$ is a set of relations specifying that all of the elements of $\mathcal{S}_i$ commute with all of the elements of $\mathcal{S}_j$. But one can easily see that this is exactly the group $\overline{G}(\Delta, \mathcal{G})$, where $\mathcal{G} = \{G_1, \ldots, G_n\}$ and Δ is the complete graph on n vertices.

6.4.2 Finite Free Products

Let $\mathcal{G} = \{G_i \mid i \in I_n\}$ and let Δ be the edgeless graph on n vertices $\{v_i \mid i \in I_n\}$. That is, $E(\Delta) = \varnothing$. Defining the free product $G_1 * \cdots * G_n$ is a bit trickier than defining the direct product $G_1 \times \cdots \times G_n$. There are several recognized approaches for defining the free product of groups.

Perhaps the simplest way to define the free product of groups is in terms of group presentations. This is the approach we used in Chapter 5 when we proved that the finite free product of totally reflected groups is totally reflected. If each factor group is given by a presentation, say $G_i = \langle \mathcal{S}_i \mid \mathcal{R}_i \rangle$, then we can define the free product in terms of the presentation shown here:

$$G_1 * \cdots * G_n = \Big\langle \bigcup_{i \in I_n} \mathcal{S}_i \Bigm| \bigcup_{i \in I_n} \mathcal{R}_i \Big\rangle$$

We can easily observe that this group presentation is exactly the defining presentation for the group $\overline{G}(\Delta, \mathcal{G})$ in this scenario.

Another approach for defining the free product of groups involves a universal mapping property. With this approach, the free product of groups is *defined* to be any group satisfying this particular universal mapping property. This is how the free product of groups is defined by Massey in [7]. If $\mathcal{G}$ and Δ are as defined above, the universal mapping property used by Massey to define $G_1 * \cdots * G_n$ corresponds to what we are referring to as *the universal mapping property determined by* $\mathcal{G}$ *and* Δ. Then according to Proposition 6.5, $G_1 * \cdots * G_n$ defined in this way must be isomorphic to the right-angled product determined by Δ and $\mathcal{G}$.

There is (at least) one other approach for defining the free product of groups, though this approach is not prevalent in the literature. We could define the free product of the groups in set $\mathcal{G} = \{G_1, \ldots, G_n\}$ using the same approach that we used at the beginning of this chapter to define the right-angled product. That is, we could define the free product of the groups $G_1, \ldots, G_n$ in terms of equivalence classes of words. Again, Δ in this situation would be the edgeless graph on the vertices $\{v_i \mid i \in I_n\}$. There would not, however, be any moves of type T in this scenario.

6.4.3 Right-Angled Coxeter and Artin Groups

For any finite simplicial graph Δ with vertices $v_1, \ldots, v_n$, we can define the right-angled Artin group determined by Δ by the following presentation:

$$A(\Delta) := \langle v_1, \ldots, v_n \mid v_i v_j = v_j v_i \iff \{v_i, v_j\} \in E(\Delta) \rangle$$

For each $i \in I_n$, suppose we take G_i to be the infinite cyclic group generated by v_i, that is, $G_i = \langle \mathcal{S}_i \mid \mathcal{R}_i \rangle$ where $\mathcal{S}_i = \{v_i\}$ and $\mathcal{R}_i = \varnothing$. Let $\mathcal{G} = \{G_1, \ldots, G_n\}$. Then the following relationship is easily observed:

$$A(\Delta) \cong \overline{G}(\Delta, \mathcal{G})$$

Suppose that in the above presentation for $A(\Delta)$ we add the relation that says $v_i = v_i^{-1}$ for each $i \in I_n$, meaning that each generator has order 2 in the group. The resulting group is known as the right-angled Coxeter group determined by Δ, which can be denoted as $W(\Delta)$. If for each $i \in I_n$ we let

$H_i = \langle \mathcal{S}_i \,\big|\, \mathcal{R}_i \rangle$, where $\mathcal{S}_i = \{v_i\}$ and $\mathcal{R}_i = \{v_i^2\}$, then we can easily observe that $W(\Delta) \cong \overline{G}(\Delta, \mathcal{G})$ where $\mathcal{G} = \{H_1, \ldots, H_n\}$.

6.5 Right-Angled Products of Totally Reflected Groups

In Corollaries 5.14 and 5.20 of the preceding chapter, we proved that any finite direct or free product of totally reflected groups is totally reflected. The principal theorem of this chapter, Theorem 6.9, will provide a generalization of these theorems to the case of the right-angled product. First, we must set the stage with the tools and terminology we will need in order to prove Lemma 6.8 and subsequently Theorem 6.9.

Since $\overline{G} = \overline{G}(\Delta, \mathcal{G})$ is generated by $\mathcal{S}_{\mathcal{G}} = \bigcup_{i \in I_n} \mathcal{S}_i$, each element of $\overline{G}$ can be written as a word in $\mathcal{S}_{\mathcal{G}}^{\pm}$, though this representation as a word need not be unique. Suppose w is a word in $\mathcal{S}_{\mathcal{G}}^{\pm}$, say $w = x_1 x_2 \cdots x_m$. Fix $i \in I_n$. An **$\mathcal{S}_i$–syllable of w** is any subword in w of the form $x_j x_{j+1} \cdots x_{j+m_j}$, where $j \in I_n$, $m_j \in \overline{I}_{m-j}$, the letters x_j, x_{j+1}, $\ldots$, x_{j+m_j} are all elements of $\mathcal{S}_i^{\pm}$, $x_{j-1} \notin \mathcal{S}_i^{\pm}$ in the event that $j \geq 2$, and $x_{j+m_j+1} \notin \mathcal{S}_i^{\pm}$ in the event that $j + m_j \leq m - 1$. As we move from left to right in the expression $w = x_1 \cdots x_m$, we can derive a unique sequence of consecutive, disjoint subwords $(w_1, \ldots, w_r)$ such that the following conditions are satisfied:

- for each $k \in I_r$, w_k is an $\mathcal{S}_{f(k)}$–syllable for some $f(k) \in I_n$
- in the event that $r \geq 2$, $f(k) \neq f(k+1)$ for all $k \in I_{r-1}$
- $w = w_1 \cdots w_r$

The sequence $(w_1, \ldots, w_r)$ will be called the **disjoint syllable decomposition** of the word w. The positive integer r will be called the **disjoint syllable length** of w and will be denoted as $l(w)$.

Suppose $w = x_1 \cdots x_m$ and $\widetilde{w} = y_1 \cdots y_{m'}$ are words in $\mathcal{S}_{\mathcal{G}}^{\pm}$ with disjoint syllable decompositions $(w_1, \ldots, w_r)$ and $(\widetilde{w}_1, \ldots, \widetilde{w}_{r'})$, respectively. Consider the concatenation $w\widetilde{w} = x_1 \cdots x_m y_1 \cdots y_{m'}$. Observe the following:

$$l(w\widetilde{w}) = l(w) + l(\widetilde{w}) \iff w_r, \widetilde{w}_1 \text{ are not both } \mathcal{S}_j\text{–syllables for any } j \in I_n$$

If w_r and $\widetilde{w}_1$ are both $\mathcal{S}_j$–syllables for some $j \in I_n$, then $l(w\widetilde{w}) = l(w)+l(\widetilde{w})-1$.

For any word w in $\mathcal{S}_{\mathcal{G}}^{\pm}$, define the **support** of w, denoted by $\operatorname{supp}(w)$, as shown here:

$$\operatorname{supp}(w) := \{v_j \in V(\Delta) \mid w \text{ contains an } \mathcal{S}_j\text{–syllable}\}$$

We can easily observe that for any words w and $\widetilde{w}$ in $\mathcal{S}_{\mathcal{G}}^{\pm}$, $\operatorname{supp}(w\widetilde{w}) = \operatorname{supp}(w) \cup \operatorname{supp}(\widetilde{w})$. Also, if w' is an $\mathcal{S}_i$–syllable of a word w in $\mathcal{S}_{\mathcal{G}}^{\pm}$, then $\operatorname{supp}(w') = \{v_i\}$.

For any $i \in I_n$, define the **link of a vertex $\boldsymbol{v_i}$** as follows:

$$\operatorname{link}(v_i) := \{v_j \in V(\Delta) \mid j \in I_n - \{i\} \text{ and } \{v_i, v_j\} \in E(\Delta)\}$$

Observe that for $i, j \in I_n$, $i \neq j$, $v_i \in \operatorname{link}(v_j)$ if and only if $v_j \in \operatorname{link}(v_i)$. Suppose that w is a word in $\mathcal{S}_{\mathcal{G}}^{\pm}$ and that $\operatorname{supp}(w) \subseteq \operatorname{link}(v_i)$ for some $i \in I_n$. Then necessarily $v_i \notin \operatorname{supp}(w)$, since $v_i \notin \operatorname{link}(v_i)$. Recall that if $s \in \mathcal{S}_i$ and $t \in \mathcal{S}_j$ for some $i, j \in I_n$, $i \neq j$, with $\{v_i, v_j\} \in E(\Delta)$, then $sts^{-1}t^{-1}$ is a relator in $\overline{G}$. Therefore, $st \overset{\overline{G}}{=} ts$. Consequently, if $\text{supp}(w) \subseteq \text{link}(v_i)$ and if $g \in G_i$, then $wg \overset{\overline{G}}{=} gw$.

Suppose w is a word in $\mathcal{S}_{\mathcal{G}}^{\pm}$ with disjoint syllable decomposition $(w_1, \ldots, w_r)$ where $r \geq 3$. Consider a subword of w of the form $w_j w_{j+1} \cdots w_{k-1} w_k$ such that the following conditions hold:

- $k - j \geq 2$
- w_j and w_k are both $\mathcal{S}_i$–syllables for some $i \in I_n$
- $\text{supp}(w_{j+1} \cdots w_{k-1}) \subseteq \text{link}(v_i)$

We will call such a subword a **reducible segment** of w. Since $\text{supp}(w_{j+1} \cdots w_{k-1}) \subseteq \text{link}(v_i)$ and since w_j and w_k are both $\mathcal{S}_i$–syllables, then $w_j w_{j+1} \cdots w_{k-1} w_k \overset{\overline{G}}{=} w_j w_k w_{j+1} \cdots w_{k-1}$. We can also see that $l(w_j w_{j+1} \cdots w_{k-1} w_k) = k - j + 1$, whereas $l(w_j w_k w_{j+1} \cdots w_{k-1}) = k - j$. This reduction of disjoint syllable length is the reason for calling $w_j w_{j+1} \cdots w_{k-1} w_k$ a reducible segment. We will say that a word w in $\mathcal{S}_{\mathcal{G}}^{\pm}$ is **graphically reduced** if it contains no reducible segments.

The concepts of graphically reduced words, support, link, and disjoint syllable length will be essential to the proof of Lemma 6.8 which follows. This lemma does the bulk of the hard work necessary to prove that a right-angled product of totally reflected groups is totally reflected.

Lemma 6.8. *Fix an arbitrary index $i \in I_n$. Any color-preserving graph reflection on $\Gamma(G_i, \mathcal{S}_i)$ which inverts the edge $e(1_{G_i}, k)$, where $k \in \mathcal{S}_i$, can be extended to a color-preserving graph reflection on $\Gamma(\overline{G}(\Delta, \mathcal{G}), \mathcal{S}_{\mathcal{G}})$ which inverts the edge $e(1_{\overline{G}(\Delta,\mathcal{G})}, k)$.*

Proof. For convenience, we will abbreviate $\Gamma(G_i, \mathcal{S}_i)$ and $\Gamma(\overline{G}(\Delta, \mathcal{G}), \mathcal{S}_{\mathcal{G}})$ as $\Gamma(G_i)$ and $\Gamma(\overline{G})$, respectively. Additionally, we will abbreviate $\overline{G}(\Delta, \mathcal{G})$ as $\overline{G}$ and $1_{\overline{G}(\Delta,\mathcal{G})}$ as $1_{\overline{G}}$.

Any color-preserving reflection on $\Gamma(G_i)$ which inverts the edge $e(1_{G_i}, k)$ can be expressed in the form $L_k\alpha$, where $\alpha \in \mathrm{Aut}(G_i, \mathcal{S}_i^{\pm})$ satisfies $\alpha(k) = k^{-1}$ and $\alpha^2 = \mathrm{id}_{G_i}$. We can extend α to a function α^* on all of $\overline{G}$ by defining $\alpha^*(s) = \alpha(s)$ for all $s \in \mathcal{S}_i$; by defining $\alpha^*(s) = s$ for all $s \in \mathcal{S}_{\mathcal{G}} - \mathcal{S}_i$; and by extending this definition linearly so that it applies to all elements of $\overline{G}$.

We must now show that $\alpha^* : \overline{G} \to \overline{G}$ is a homomorphism. We will do this by showing that α^* respects the relators of $\overline{G}$. First, suppose $r \in \mathcal{R}_i$. Since $\alpha : G_i \to G_i$ is an automorphism, $\alpha(r) \overset{G_i}{=} 1_{G_i}$. However, as an element of $\overline{G}$, $1_{G_i} = 1_{\overline{G}}$, and so we can see that $\alpha^*(r) = \alpha(r) = 1_{\overline{G}}$. Next, suppose $r \in \mathcal{R}_j$ for some $j \neq i$, meaning that $r \overset{G_j}{=} 1_{G_j} \overset{\overline{G}}{=} 1_{\overline{G}}$. Then $\alpha^*(r) = r \overset{G_j}{=} = 1_{\overline{G}}$. We have shown, then, that α^* respects the relators coming from $\bigcup_{j \in I_n} \mathcal{R}_j$. Let us now consider what happens when $r \in \mathcal{R}_{jk}$, where $j, k \in I_n$, $j \neq k$, and $\{v_j, v_k\} \in E(\Delta)$. Then for some $s \in \mathcal{S}_j$, $t \in \mathcal{S}_k$, we have $r = sts^{-1}t^{-1}$. If $j = i$, then $k \neq i$, and we can observe the following:

$$\alpha^*(r) = \alpha^*(sts^{-1}t^{-1}) = \alpha(s)t\alpha(s^{-1})t^{-1} = \alpha(s)t\alpha(s)^{-1}t^{-1} = 1_{\overline{G}}$$

The last equality comes from the fact that $\alpha(s)t\alpha(s)^{-1}t^{-1} \in \mathcal{R}_{jk}$, since $\alpha(s) \in \mathcal{S}_j$. If $k = i$, then $j \neq i$, and a symmetric argument to the one just given would show that $\alpha^*(r) = 1_{\overline{G}}$. So suppose now that $j \neq i$ and $k \neq i$. Then $\alpha^*(r) = \alpha^*(sts^{-1}t^{-1}) = sts^{-1}t^{-1} = 1_{\overline{G}}$. Therefore, α^* respects the relators of $\overline{G}$ and must be a homomorphism.

We can easily observe that $(\alpha^*)^2 = \mathrm{id}_{\overline{G}}$, since $\alpha^2 = \mathrm{id}_{G_i}$. So $(\alpha^*)^2(s) =$

$\alpha^2(s) = s$ if $s \in \mathcal{S}_i$ and $(\alpha^*)^2(s) = s$ if $s \in \mathcal{S}_j$ with $j \neq i$. Therefore, for any $g \in \overline{G}$, $(\alpha^*)^2(g) = g$. From this we can conclude that α^* is bijective and thus is an automorphism. Moreover, $(\alpha^*)^{-1} = \alpha^*$. We can easily see by definition of α^* that $\alpha^*\left(\mathcal{S}_{\mathcal{G}}^{\pm}\right) = \mathcal{S}_{\mathcal{G}}^{\pm}$. Thus, $\alpha^* \in \operatorname{Aut}(\overline{G}, \mathcal{S}_{\mathcal{G}}^{\pm})$.

By Theorem 3.8, we can see that $L_k\alpha^*$ is a color-preserving graph automorphism on $\Gamma(\overline{G})$. Note that $L_k\alpha^*$ inverts the edge $e(1_{\overline{G}}, k)$ in $\Gamma(\overline{G})$, since $L_k\alpha^*(1_{\overline{G}}) = k\alpha^*(1_{\overline{G}}) = k1_{\overline{G}} = k$ and $L_k\alpha^*(k) = k\alpha(k) = kk^{-1} = 1_{\overline{G}}$. Our goal now is to show that $L_k\alpha^*$ is a reflection on $\Gamma(\overline{G})$.

First, we must show that $L_k\alpha^*$ has order 2. Suppose that $g \in \overline{G}$. Then

$$\begin{aligned}(L_k\alpha^*)^2(g) &= L_k\alpha^*(L_k\alpha^*(g)) = k\alpha^*(k\alpha^*(g)) \\ &= k\alpha^*(k)(\alpha^*)^2(g) = k\alpha(k)g = kk^{-1}g = g\end{aligned}$$

which illustrates that $(L_k\alpha^*)^2$ fixes all of the vertices of $\Gamma(\overline{G})$ and thus must be equal to $\mathrm{id}_{\Gamma(\overline{G})}$.

Next, consider the set of edges defined below:

$$\begin{aligned}\mathcal{E} := \{e(xg, b) \,\big|\, g \in G_i,\ b \in \mathcal{S}_i^{\pm},\ e(g,b) &\text{ is inverted by } L_k\alpha \text{ in } \Gamma(G_i), \\ x \in \overline{G},\ \operatorname{supp}(x) \subseteq \operatorname{link}(v_i)\}&\end{aligned}$$

If we take $g = 1_{G_i}$, $b = k$, and $x = 1_{\overline{G}}$ as in the definition of $\mathcal{E}$, then we can see that the edge $e(1_{\overline{G}}1_{G_i}, k) = e(1_{\overline{G}}, k)$ in $\Gamma(\overline{G})$ must be an edge in set $\mathcal{E}$, since the edge $e(1_{G_i}, k)$ is inverted by $L_k\alpha$ in $\Gamma(G_i)$.

We claim that all of the edges in $\mathcal{E}$ are inverted by $L_k\alpha^*$. To see this, assume that x, g, b are as in the definition of $\mathcal{E}$. Observe the following properties:

- $\alpha^*(x) = x$, since $\text{supp}(x) \subseteq \text{link}(v_i)$
- $xk = kx$, since $k \in G_i$ and $\text{supp}(x) \subseteq \text{link}(v_i)$
- $\alpha^*(g) = \alpha(g)$, since $g \in G_i$
- $k\alpha(g) = gb$, since the edge $e(g, b)$ gets inverted by $L_k\alpha$ in $\Gamma(G_i)$
- $k\alpha(gb) = g$, since the edge $e(g, b)$ gets inverted by $L_k\alpha$ in $\Gamma(G_i)$

From these properties we may deduce the following:

$$L_k\alpha^*(xg) = k\alpha^*(x)\alpha^*(g) = kx\alpha(g) = xk\alpha(g) = xgb$$
$$L_k\alpha^*(xgb) = k\alpha^*(x)\alpha^*(gb) = kx\alpha(gb) = kx\alpha(gb) = xk\alpha(gb) = xg$$

Therefore, every edge in $\mathcal{E}$ is inverted by $L_k\alpha^*$, as claimed.

Our next claim is that the set $\mathcal{E}$ separates $\Gamma(\overline{G})$. In order to prove this claim, we will suppose, instead, that $\mathcal{E}$ does *not* separate $\Gamma(\overline{G})$. Then there must exist at least one path in $\Gamma(\overline{G})$ from $1_{\overline{G}}$ to k which does not contain an edge from $\mathcal{E}$, since $e(1_{\overline{G}}, k)$ is inverted by $L_k\alpha^*$. Consequently, the set $\mathcal{P} = \{$all paths in $\Gamma(\overline{G})$ from $1_{\overline{G}}$ to k which contain no edge from $\mathcal{E}\}$ must be nonempty. Choose a path in $\mathcal{P}$, say γ, that is of shortest disjoint syllable length. Let $w = w(\gamma)$ be the word in $\mathcal{S}_{\mathcal{G}}$ which is associated with path γ.

Suppose that w is graphically reduced. Then $l(w) = 1$, since w is a word which is equivalent to $k \in \mathcal{S}_i^{\pm}$. Therefore, w is, in fact, a word in $\mathcal{S}_i^{\pm}$, meaning that γ can also be thought of as a path in $\Gamma(G_i)$ from 1_{G_i} to k. Since $L_k\alpha$ is a reflection on $\Gamma(G_i)$ which inverts the edge $e(1_{G_i}, k)$, we must have that the path γ in $\Gamma(G_i)$ contains an edge which is inverted by $L_k\alpha$. Therefore, there exists an edge $e(g,b)$ in γ, with $g \in G_i$ and $b \in \mathcal{S}_i^{\pm}$, such that $e(g,b)$ is inverted by $L_k\alpha$ acting on $\Gamma(G_i)$. But this edge is clearly in $\mathcal{E}$, if we take $x = 1_{\overline{G}}$. So γ does contain an edge from $\mathcal{E}$, which contradicts our choice of $\gamma \in \mathcal{P}$. Therefore, w cannot be graphically reduced.

Since w is not graphically reduced, it must contain at least one reducible segment. Consider the earliest occurring reducible segment in w, say $z_1 y z_2$, where z_1 and z_2 are both words in $\mathcal{S}_j^{\pm}$ for some $j \in I_n$ and $\operatorname{supp}(y) \subseteq \operatorname{link}(v_j)$. Then we can write w in terms of disjoint subwords as $w = u_1 z_1 y z_2 u_2$, where in addition to the conditions just mentioned we have the following:

- $|w| = |u_1 z_1 y z_2 u_2| = |u_1| + |z_1| + |y| + |z_2| + |u_2|$
- $u_1 z_1 y$ is graphically reduced

Let γ_1, γ_2, γ_3, γ_4, and γ_5 be the consecutive subpaths of γ corresponding to the subwords mentioned above, that is,

$$w(\gamma_1) = u_1,\ w(\gamma_2) = z_1,\ w(\gamma_3) = y,\ w(\gamma_4) = z_2,\ w(\gamma_5) = u_2$$

Consider the new path $\widetilde{\gamma}$ in $\Gamma(G)$ with $\widetilde{w} = w(\widetilde{\gamma}) = u_1 z_2 z_2 y u_2$. Let $\widetilde{\gamma}_1$, $\widetilde{\gamma}_2$, $\widetilde{\gamma}_3$, $\widetilde{\gamma}_4$, and $\widetilde{\gamma}_5$ be the consecutive subpaths of $\widetilde{\gamma}$ corresponding to the subwords of $\widetilde{w}$ as shown here:

$$w(\widetilde{\gamma}_1) = u_1,\ w(\widetilde{\gamma}_2) = z_1,\ w(\widetilde{\gamma}_3) = z_2,\ w(\widetilde{\gamma}_4) = y,\ w(\widetilde{\gamma}_5) = u_2$$

Figure 6.10 illustrates the relationship between γ and $\widetilde{\gamma}$.

It is clear that $\widetilde{\gamma}$ has shorter length than γ, since z_1 and z_2 come from the same factor group. Therefore, $\widetilde{\gamma}$ must contain an edge in $e \in \mathcal{E}$, since we

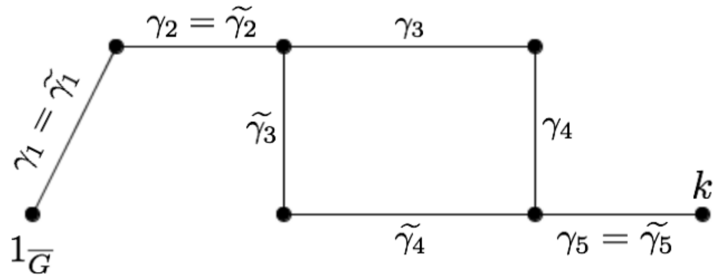

Figure 6.10: Visualizing the effect of graphical reduction on paths

are assuming that γ is the shortest path from $1_{\overline{G}}$ to k in $\Gamma(\overline{G})$ which does not contain an edge in $\mathcal{E}$. As a member of $\mathcal{E}$, the edge e must be of the form $e(xg, b)$, where $e(g, b)$ is an edge in $\Gamma(G_i)$ which is inverted by $L_k\alpha$, $x \in \overline{G}$, and $\operatorname{supp}(x) \subseteq \operatorname{link}(v_i)$. Clearly, the edge e in $\widetilde{\gamma}$ cannot be on any of the subpaths $\widetilde{\gamma}_1$, $\widetilde{\gamma}_2$ or $\widetilde{\gamma}_5$ of $\widetilde{\gamma}$, since this would imply that e is an edge in γ as well. Thus, the edge e is either on the $\widetilde{\gamma}_3$ or $\widetilde{\gamma}_4$ subpath of $\widetilde{\gamma}$.

Case i: Assume that e is on the $\widetilde{\gamma}_4$–subpath of $\widetilde{\gamma}$. Since $e = e(xg, b)$, it follows that we can write y as hbh', where $h, h' \in \overline{G}$. Since $\operatorname{supp}(x) \subseteq \operatorname{link}(v_i)$, we must have $x = u_1 z_1 z_2 h_1$ and $g = h_2$, where h_1 and h_2 are consecutive, disjoint subwords of h and where $h_1 \notin G_i - \{1\}$ but $h_2 \in G_i$.

Consider the edge $e(\widetilde{x}g, b)$, where $\widetilde{x} = u_1 z_1 h_1$. As we know from before, $e(g, b)$ is an edge in $\Gamma(G_i)$ which is inverted by $L_k\alpha$. Also, $\widetilde{x} \in \overline{G}$ with $\operatorname{supp}(\widetilde{x}) \subseteq \operatorname{supp}(x) \subseteq \operatorname{link}(v_i)$. This shows that the edge $e(\widetilde{x}g, b)$ is in the set $\mathcal{E}$. However, this edge $e(\widetilde{x}g, b)$ is an edge in the γ_3–subpath of γ, implying that γ *does* contain an edge from $\mathcal{E}$. But this is a contradiction! We assumed that γ contained no edge from $\mathcal{E}$. Therefore, this case is not actually possible.

Case ii: Assume that e is on the $\widetilde{\gamma}_3$–subpath of $\widetilde{\gamma}$. Since $e = e(xg, b)$, it follows that we can write z_2 as lbl', where $l, l' \in G(\Delta, \mathcal{G})$. Since $b \in \mathcal{S}_i^{\pm}$, it must be the case that $G_j = G_i$. Furthermore, since $\operatorname{supp}(x) \subseteq \operatorname{link}(v_i)$ and since $g \in G_i$, we must have $x = u_1$ and $g = z_1 l$.

Consider the edge $e(\widetilde{x}g, b)$, where $\widetilde{x} = u_1 y$. Since z_1 and y commute, we can see that $\widetilde{x}g = u_1 y z_1 l = u_1 z_1 y l$. Also, $u_1 z_1 y l = w(\gamma')$, where γ' is the subpath of γ consisting of the subpaths γ_1, γ_2, γ_3, and an initial segment of γ_4. Therefore, the edge $e(\widetilde{x}g, b)$ is on the path γ. As we know from before, $e(g, b)$ is an edge in $\Gamma(G_i)$ which is inverted by $L_k\alpha$. Notice that $\mathrm{supp}(\widetilde{x}) \subseteq \mathrm{link}(v_i)$, since $\mathrm{supp}(x) = \mathrm{supp}(u_1) \subseteq \mathrm{link}(v_i)$ and $\mathrm{supp}(y) \subseteq \mathrm{link}(v_i)$. We have now shown that the edge $e(\widetilde{x}g, b)$ is in the set $\mathcal{E}$ and is on the path γ. But this is a contradiction! We assumed that γ contained no edge from $\mathcal{E}$. Therefore, this case is not actually possible.

We have shown then that our assumption that w is not graphically reduced leads only to impossibilities, and so w must be graphically reduced. But we already showed this cannot happen! Therefore, our initial assumption that implied that $\mathcal{P} \neq \varnothing$ must be false. But $\mathcal{P} = \varnothing$ implies that all paths in $\Gamma(\overline{G})$ from $1_{\overline{G}}$ to k *must* contain an edge from $\mathcal{E}$.

Therefore, the set $\mathcal{E}$ separates the graph $\Gamma(\overline{G})$.

To summarize we started with a color-preserving graph reflection, $L_k\alpha$, which inverts the edge $e(1_{G_i}, k)$ in $\Gamma(G_i)$. We extended $L_k\alpha$ in a natural way to create $L_k\alpha^*$, a color-preserving graph automorphism acting on $\Gamma(\overline{G})$. Next, we showed that $(L_k\alpha^*)^2 = \mathrm{id}_{\Gamma(\overline{G})}$. We then defined a set $\mathcal{E}$ consisting of edges in $\Gamma(\overline{G})$ which are all inverted by $L_k\alpha^*$ and which separate the graph $\Gamma(\overline{G})$. We also showed that the edge $e(1_{\overline{G}}, k)$ is in $\mathcal{E}$. Therefore, we have shown that $L_k\alpha^*$ is a color-preserving reflection on $\Gamma(\overline{G}, \mathcal{S}_{\mathcal{G}})$ that inverts the edge $e(1_{\overline{G}}, k)$, which is what we set out to prove.

$\square$

We now are ready to state and prove the paramount theorem of this chapter.

Theorem 6.9. *If* $(G_i, \mathcal{S}_i)$ *is a totally reflected system for each* $i \in I_n$*, then* $\big(\overline{G}(\Delta, \mathcal{G}), \mathcal{S}_{\mathcal{G}}\big)$ *is a totally reflected system.*

Proof. Assume that $(G_j, \mathcal{S}_j)$ is a totally reflected system for each $j \in I_n$. Let $s \in \mathcal{S}_{\mathcal{G}}$ be arbitrary. Consider the edge $e(1_{\overline{G}}, s)$ in $\Gamma(\overline{G})$. Then $s \in \mathcal{S}_i$ for some $i \in I_n$. We must show that there exists a color-preserving reflection acting on $\Gamma(\overline{G})$ which inverts the edge $e(1_{\overline{G}}, s)$.

Since the factor group G_i is totally reflected with respect to the generating set $\mathcal{S}_i$, there exists a color-preserving graph reflection acting on $\Gamma(G_i)$ of the form $L_s\alpha$ which inverts the edge $e(1_{G_i}, s)$, where $\alpha \in \mathrm{Aut}(G_i, \mathcal{S}_i^{\pm})$ and $\alpha(s) = s^{-1}$. Then by Lemma 6.8, we may extend $L_s\alpha$ to a color-preserving graph reflection acting on $\Gamma(\overline{G})$ which inverts the edge $e(1_{\overline{G}}, s)$. Since $s \in \mathcal{S}_{\mathcal{G}}$ was arbitrary, Proposition 5.5 allows us to conclude that the right-angled product system $\big(\overline{G}(\Delta, \mathcal{G}), \mathcal{S}_{\mathcal{G}}\big)$ is totally reflected.

∎

Chapter 7

Strongly Totally Reflected Groups

7.1 Definitions, Examples, and Basic Properties

We will say that a group system $(G, \mathcal{S})$ is **strongly totally reflected** (or **strongly t.r.** or **s.t.r.**) if there is a reflection group G_R acting on $\Gamma = \Gamma(G, \mathcal{S})$ such that $G_R \subseteq \text{Aut}_{\text{c.p.}}(\Gamma)$ and such that for each edge e in Γ there exists a reflection in G_R which inverts e. Additionally, we will say that the group G is **s.t.r. with respect to** $\mathcal{S}$ if the system $(G, \mathcal{S})$ is s.t.r. As we did in the case of totally reflected (t.r.) groups, we will say that a group G is **s.t.r.** (or is an **s.t.r. group**) if it has a generating set $\mathcal{S}$ for which $(G, \mathcal{S})$ is an s.t.r. system.

Example 7.1. If W is a Coxeter group with fundamental generating set $\mathcal{S}$, then the system $(W, \mathcal{S})$ is s.t.r. As we have discussed previously, the group W itself is a reflection group acting on $\Gamma = \Gamma(W, \mathcal{S})$. We know that $W = \text{Aut}_{\text{c.f.}}(\Gamma)$, by Lemma 3.5, and that $\text{Aut}_{\text{c.f.}}(\Gamma) \subseteq \text{Aut}_{\text{c.p.}}(\Gamma)$. If $e(g,s)$ is any edge in Γ, where $g \in W$ and $s \in \mathcal{S}$, then $L_g L_s (L_g)^{-1} = L_{gsg^{-1}}$ is an element of $W = \text{Aut}_{\text{c.f.}}(\Gamma)$ which inverts the edge $e(g,s)$.

◇

If $(G, \mathcal{S})$ is a t.r. group system, then we know that for any edge e in Γ there exists a color-preserving graph reflection on Γ, say r_e, which inverts the edge e. Consider the group $R := \langle r_e \mid e \in E(\Gamma) \rangle$. R is generated by reflections and is a subset of $\text{Aut}_{\text{c.p.}}(\Gamma)$, since each $r_e \in \text{Aut}_{\text{c.p.}}(\Gamma)$. For each $e \in E(\Gamma)$, the reflection r_e is vertex free, and thus edge free, on Γ. However, there may be non-generator elements of R which do not act edge freely on Γ. The next example will illustrate this.

Example 7.2. Suppose $G = \mathbb{Z}_2 * \mathbb{Z} = \langle a, b \mid a^2 = 1_G \rangle$ and $\mathcal{S} = \{a, b\}$. The Cayley graph $\Gamma = \Gamma(G, \mathcal{S})$ is a three-regular tree with an undirected a-colored edge, an incoming b-colored edge, and an outgoing b-colored edge incident to each vertex. Γ is shown in Figure 7.1.

If we think of $\mathbb{Z}_2$ as being generated by a and $\mathbb{Z}$ as being generated by b, we know from previous examples that $(\mathbb{Z}_2, \{a\})$ and $(\mathbb{Z}, \{b\})$ are t.r. systems. Then Theorem 5.19 tells us that $(\mathbb{Z}_2 * \mathbb{Z}, \mathcal{S})$ is a t.r. system also. Therefore, for any $e \in E(\Gamma)$, we know that we can find at least one color-preserving reflection inverting e.

By definition of Cayley graph, $e_0 = e(1_G, a)$ and $\overline{e_0} = e(a, a)$ are distinct edges in Γ. In Figure 7.1 we have drawn just one of the edges coming from the pair $\{e_0, \overline{e_0}\}$, as is conventional when drawing Cayley graphs. Suppose ϕ is a group automorphism of G defined by stating that $\phi(a) = a$ and $\phi(b) = b^{-1}$.

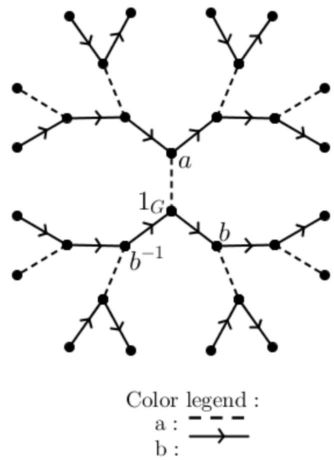

Figure 7.1: A portion of the Cayley graph $\Gamma(\mathbb{Z}_2 * \mathbb{Z}, \{a, b\})$

Note that L_a and $L_a\phi$ are distinct color-preserving graph reflections on Γ. Moreover, both of these reflections have the same wall consisting of e_0 and $\overline{e_0}$. Using the notation preceding this example, let $r_{e_0} = L_a$ and $r_{\overline{e_0}} = L_a\phi$. For any $e \in E(\Gamma) - \{e_0, \overline{e_0}\}$, let r_e be any color-preserving reflection inverting e.

Observe that r_{e_0} and $r_{\overline{e_0}}$ are both elements of $\{r_e \mid e \in E(\Gamma)\}$, and so $r_{e_0}r_{\overline{e_0}} \in R = \langle r_e \mid e \in E(\Gamma)\rangle$. We can easily observe that the composition $r_{e_0}r_{\overline{e_0}} = L_aL_a\phi = \phi$. Since ϕ fixes the elements 1_G and a in G, the composition $r_{e_0}r_{\overline{e_0}}$ must fix those points in Γ and therefore must fix the entire edge $e(1_G, a)$ in Γ. Note also that $r_{e_0}r_{\overline{e_0}}(b) = \phi(b) = b^{-1} \neq b$, and so $r_{e_0}r_{\overline{e_0}}$ does not fix the entire graph Γ. So the action of R on Γ is not edge free. Consequently, R is not a reflection group acting on Γ.

$\diamond$

The preceding example shows us that if we hope to find a reflection group G_R acting on $\Gamma(G, \mathcal{S})$ such that G_R is made up of elements of $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$, we will have to be more intentional in our choice of reflections which we use to generate G_R.

Example 7.3. Consider once again the setting from Example 7.2. Let $\widetilde{R} := \langle L_a, L_b\phi, L_{b^{-1}}\phi\rangle$, where ϕ is the same as it was in the earlier example, with $\phi(a) = a$ and $\phi(b) = b^{-1}$. Note that L_a, $L_b\phi$, and $L_{b^{-1}}\phi$ are color-preserving reflections on $\Gamma = \Gamma(G, \mathcal{S})$ which invert the edges $e(1_G, a)$, $e(1_G, b)$, and $e(1_G, b^{-1})$, respectively. We claim that $\widetilde{R}$ satisfies the following properties: the group $\widetilde{R}$ acts edge freely on Γ and thus is a reflection group on Γ; $\widetilde{R} \subseteq \mathrm{Aut}_{\text{c.p.}}(\Gamma)$ since $L_a, L_b\phi, L_{b^{-1}}\phi \in \mathrm{Aut}_{\text{c.p.}}(\Gamma)$; and for any edge e in Γ, there exists an element of $\widetilde{R}$ which inverts the edge e. (Theorem 7.5 and Proposition 7.6 will show that these claims are true.) Therefore, the existence of a group $\widetilde{R}$ with the properties just outlined allows us to conclude that $(G, \mathcal{S}) = (\mathbb{Z}_2 * \mathbb{Z}, \{a, b\})$ is s.t.r.

◇

Proposition 7.4. *Let $(G, \mathcal{S})$ be a totally reflected group system and G_R be a reflection group acting on $\Gamma = \Gamma(G, \mathcal{S})$ such that $G_R \subseteq \mathrm{Aut}_{\text{c.p.}}(\Gamma)$. If for any $s \in \mathcal{S}^{\pm}$ there exists a reflection in G_R which inverts the edge $e(1_G, s)$, then for any edge $e \in E(\Gamma)$ there exists a reflection in G_R which inverts e. Consequently, $(G, \mathcal{S})$ is strongly totally reflected.*

Proof. For each $s \in \mathcal{S}^{\pm}$, let r_s be a reflection in G_R which inverts the edge $e(1_G, s)$. Since $G_R \subseteq \mathrm{Aut}_{\text{c.p.}}(\Gamma)$ and since $r_s \in G_R$ is a reflection, we can write $r_s = L_s\phi_s$, where $\phi_s \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$ satisfies $\phi_s(s) = s^{-1}$ and $(\phi_s)^2 = \mathrm{id}_G$. Since G_R is a reflection group, r_s must be the *unique* reflection in G_R which inverts the edge $e(1_G, s)$. Consider an arbitrary edge e in Γ which is not incident to the 1_G vertex in Γ. Then there exists $g \in G - \{1_G\}$ and $t \in \mathcal{S}^{\pm}$ such that

$e = e(g,t)$. Let γ be a path with edge length k in Γ from 1_G to g. Since $g \neq 1_G$, we must have $k \geq 1$. As elements of G, $g = w(\gamma)$, the word in $\mathcal{S}^{\pm}$ which is written by the path γ. We claim that there exists a reflection in G_R which inverts the edge $e = e(g,t)$. We will prove this claim by inducting on k.

If $k = 1$, then $g = s_1$ for some $s_1 \in \mathcal{S}^{\pm}$ and $e = e(s_1, t)$. Consider the element of G_R which is given by $r = r_{s_1} r_{t'} r_{s_1}$, where $t' = \phi_{s_1}(t) \in \mathcal{S}^{\pm}$. Since $r_{t'}$ is a reflection, r is also a reflection since it is a conjugate of a reflection. Note that the reflection $r_{t'}$ inverts the edge $e(1_G, t')$ in Γ and $e(1_G, t') = e(1_G, \phi_{s_1}(t)) = L_{s_1}\phi_{s_1}\big(e(s_1,t)\big) = r_{s_1}\big(e(g,t)\big)$. Then we can observe the following:

$$r_{s_1} r_{t'} r_{s_1}\big(e(g,t)\big) = r_{s_1}\Big(\overline{r_{s_1}(e(g,t))}\Big) = \overline{r_{s_1} r_{s_1}(e(g,t))} = \overline{e(g,t)}$$

Therefore, r inverts the edge $e = e(g,t)$. Therefore, the claim is true when $k = 1$.

Now suppose the claim is true for some k. That is, suppose that when the path in Γ from 1_G to h has length k, then the edge with initial vertex h can be inverted by some reflection in G_R. If γ has length $k+1$, then $g = w(\gamma) = s_1 s_2 \cdots s_k s_{k+1}$ where $s_i \in \mathcal{S}^{\pm}$ for each $i \in I_{k+1}$.Therefore, $\gamma = (e_1, ..., e_{k+1})$, where $e_1 = e(1_G, s_1)$ and $e_{i+1} = e(s_1 s_2 \cdots s_i, s_{i+1})$ for each $i \in I_k$. Consider the path $\widetilde{\gamma}$ in Γ given by $(r_{s_1}(e_2), r_{s_1}(e_3), ..., r_{s_1}(e_{k+1}))$. Note that $w(\widetilde{\gamma}) = \phi_{s_1}(s_2)\phi_{s_1}(s_3) \cdots \phi_{s_1}(s_{k_1}) = \phi_{s_1}(s_2 s_3 \cdots s_{k+1})$. Since r_{s_1} inverts the edge $e(1_G, s_1)$, we must have that $\iota(\widetilde{\gamma}) = \iota(r_{s_1}(e_2)) = r_{s_1}(\iota(e_2)) = r_{s_1}(s_1) = 1_G$. Therefore, $\widetilde{\gamma}$ is a path in Γ beginning at 1_G which has edge length of k. So by the inductive hypothesis, the edge $\widetilde{e} = e(\phi_{s_1}(s_2 s_3 \cdots s_{k+1}), \phi_{s_1}(t))$ is inverted by some reflection, say $\widetilde{r}$, in G_R. Note the following:

$$
\begin{aligned}
\widetilde{e} &= e(\phi_{s_1}(s_2 s_3 \cdots s_{k+1}), \phi_{s_1}(t)) \\
&= e(1_G \phi_{s_1}(s_2 s_3 \cdots s_{k+1}), \phi_{s_1}(t)) \\
&= e(r_{s_1}(s_1)\phi_{s_1}(s_2 s_3 \cdots s_{k+1}), \phi_{s_1}(t)) \\
&= e(L_{s_1}\phi_{s_1}(s_1 s_2 s_3 \cdots s_{k+1}), \phi_{s_1}(t)) \\
&= L_{s_1}\big(e(\phi_{s_1}(g), \phi_{s_1}(t))\big) \\
&= L_{s_1}\phi_{s_1}\big(e(g,t)\big) \\
&= r_1(e(g,t))
\end{aligned}
$$

Since $\widetilde{r}$ and r_{s_1} are reflections in G_R, $r_{s_1}\widetilde{r}r_{s_1}$ is also a reflection in G_R. Observe the following:

$$
\begin{aligned}
r_{s_1}\widetilde{r}r_{s_1}\big(e(g,t)\big) &= r_{s_1}\widetilde{r}\big(r_{s_1}(e(g,t))\big) = r_{s_1}\widetilde{r}(\widetilde{e}) = r_{s_1}\big(\overline{\widetilde{e}}\big) \\
&= \overline{r_{s_1}(\widetilde{e})} = \overline{r_{s_1}(r_{s_1}(e(g,t)))} = \overline{e(g,t)}
\end{aligned}
$$

Therefore, $r_{s_1}\widetilde{r}r_{s_1}$ is a reflection in G_R which inverts the edge $e(g,t)$. This completes the proof by induction. Therefore, we have now shown the following: if for any $s \in \mathcal{S}^{\pm}$ there exists a reflection in G_R which inverts the edge $e(1_G, s)$, then for *any* edge $e \in E(\Gamma)$ there exists a reflection in G_R which inverts e. Along with our fundamental assumptions, this allows us to conclude that $(G, \mathcal{S})$ is strongly totally reflected.

□

7.2 A Sufficient Condition for a T.R. Group to be S.T.R.

Theorem 7.5. *Let $(G, \mathcal{S})$ be a totally reflected group system. For each $s \in \mathcal{S}^{\pm}$, let $L_s\phi_s$ be a color-preserving reflection on $\Gamma = \Gamma(G, \mathcal{S})$ inverting the edge $e(1_G, s)$. Assume that $\phi_s = \phi_{s^{-1}}$ for all $s \in \mathcal{S}$. If Condition $(\star\star)$ holds for all (nonempty) words $s_{i_1}s_{i_2}\cdots s_{i_k}$ in $\mathcal{S}^{\pm}$, then $(G, \mathcal{S})$ is strongly totally reflected.*

$$(\star\star): \qquad s_{i_1} \cdot [\phi_{s_{i_1}}(s_{i_2})] \cdot [\phi_{s_{i_1}}\phi_{s_{i_2}}(s_{i_3})] \cdots [\phi_{s_{i_1}} \cdots \phi_{s_{i_{k-1}}}(s_{i_k})] = 1_G$$

$$\Downarrow$$

$$\phi_{s_{i_1}} \cdots \phi_{s_{i_k}} = id_\Gamma$$

Proof. Let $G_R := \langle L_s\phi_s \mid s \in \mathcal{S}^{\pm} \rangle$. Note that $G_R \subseteq \mathrm{Aut}_{\mathrm{c.p.}}(\Gamma)$, since $L_s\phi_s \in \mathrm{Aut}_{\mathrm{c.p.}}(\Gamma)$ for all $s \in \mathcal{S}^{\pm}$. If we can show that the action of G_R is edge free on Γ, then Proposition 7.4 will allow us to conclude that $(G, \mathcal{S})$ is s.t.r. Assume that Condition $(\star\star)$ holds for all (nonempty) words $s_{i_1}s_{i_2}\cdots s_{i_k}$ in $\mathcal{S}^{\pm}$.

Suppose that $L_g\phi \in \mathrm{Aut}_{\mathrm{c.p.}}(\Gamma)$, where $g \in G$ and $\phi \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$. Also, suppose that $L_g\phi$ fixes an edge in Γ of the form $e(1_G, s)$ for some $s \in \mathcal{S}^{\pm}$. In particular, $L_g\phi$ must fix both endpoints of $e(1_G, s)$ and so we must have that $g = 1_G$ and $\phi(s) = s$. Therefore, $L_g\phi = \phi$, which shows us that the only elements of $\mathrm{Aut}_{\mathrm{c.p.}}(\Gamma)$ which can fix an edge incident to the identity vertex in Γ must come specifically from $\mathrm{Aut}(G, \mathcal{S}^{\pm})$.

We will show now that $G_R \cap \mathrm{Aut}(G, \mathcal{S}^{\pm}) = \{\mathrm{id}_\Gamma\}$. To see this, suppose that $\phi \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$ and $\phi \in G_R$. Then ϕ can be written as $\phi = L_{t'_1}\phi_{t'_1} \cdots L_{t'_m}\phi_{t'_m} = L_{t'_1}\phi_{t_1} \cdots L_{t'_m}\phi_{t_m}$, where $t_i \in \mathcal{S}$ and $t'_i = t_i$ or t_i^{-1} for each $i \in I_m$. The expression for ϕ can be rewritten as shown here:

$$\begin{aligned}\phi &= L_{t'_1}\phi_{t_1} \cdots L_{t'_m}\phi_{t_m} \\ &= L_{t'_1\phi_{t_1}(t'_2)\phi_{t_1}\phi_{t_2}(t'_3)\cdots\cdots\phi_{t_1}\cdots\phi_{t_{m-1}}(t'_m)}\phi_{t_1} \cdots \phi_{t_m}\end{aligned}$$

But in order for ϕ to equal the last expression, the following must be true:

$$t'_1\phi_{t_1}(t'_2)\phi_{t_1}\phi_{t_2}(t'_3) \cdots \phi_{t_1} \cdots \phi_{t_{m-1}}(t'_m) = 1_G$$

Condition $(\star\star)$ then gives us that $\phi_{t_1}\phi_{t_2}\cdots\phi_{t_m} = \mathrm{id}_\Gamma$. Therefore, we now have that $\phi = L_{1_G}\phi_{t_1}\cdots\phi_{t_m} = \mathrm{id}_\Gamma$, which allows us to conclude that $G_R \cap \mathrm{Aut}(G, \mathcal{S}^\pm) = \{\mathrm{id}_\Gamma\}$.

We must show now that an arbitrary edge $e(g,t)$ in Γ, where $g \in G$ and $t \in \mathcal{S}^\pm$, cannot be fixed by any non-identity element of G_R. Suppose that $\beta \in G_R$ fixes the edge $e(g,t)$. We claim that there exists $\phi \in \mathrm{Aut}(G,\mathcal{S}^\pm)$ such that $L_g\phi \in G_R$. To see this, we will begin by writing g as a word in $\mathcal{S}^\pm$, say $g = u_1\cdots u_m$. Let $\epsilon_1 = 1$ and for integers $i \geq 2$ define ϵ_i as shown here:

$$\epsilon_i := \begin{cases} 1, & \text{if } \exists \text{ an even number of indices } j \in I_{i-1} \text{ such that } u_j = u_i \text{ or } u_i^{-1} \\ -1, & \text{if } \exists \text{ an odd number of indices } j \in I_{i-1} \text{ such that } u_j = u_i \text{ or } u_i^{-1} \end{cases}$$

For each $i \in I_m$, let $w_i := u_i^{\epsilon_i}$. Then $w_1 = u_1$, and for each $i \in I_m$ with $i \geq 2$ we have $\phi_{w_1}\phi_{w_2}\cdots\phi_{w_{i-1}}(w_i) = u_i$. Note that $L_{w_i}\phi_{w_i} \in G_R$ for each $i \in I_m$. Then $L_{w_1}\phi_{w_1}L_{w_2}\phi_{w_2} \cdots L_{w_m}\phi_{w_m} \in G_R$. Let $x \in G$ be arbitrary and observe the following:

$$\begin{aligned} L_{w_1}\phi_{w_1}\cdots L_{w_m}\phi_{w_m}(x) &= w_1\phi_{w_1}(w_2)\cdots\phi_{w_1}\phi_{w_2}\cdots\phi_{w_{m-1}}(w_m)\phi_{w_1}\cdots\phi_{w_m}(x) \\ &= u_1u_2\cdots u_m\phi_{w_1}\cdots\phi_{w_m}(x) \\ &= g\phi_{w_1}\cdots\phi_{w_m}(x) \\ &= L_g\phi_{w_1}\cdots\phi_{w_m}(x) \end{aligned}$$

Therefore, if we take $\phi := \phi_{w_1}\phi_{w_2}\cdots\phi_{w_m}$, then $\phi \in \mathrm{Aut}(G,\mathcal{S}^{\pm})$ and $L_g\phi = L_g\phi_{w_1}\cdots\phi_{w_m} = L_{w_1}\phi_{w_1}\cdots L_{w_m}\phi_{w_m} \in G_R$, as desired. Since $L_g\phi$, $\beta \in G_R$, we also have $(L_g\phi)^{-1}\beta L_g\phi \in G_R$.

Since β fixes the edge $e(g,t)$, we must have $\beta(g) = g$ and $\beta(gt) = gt$. Since $L_g\phi(1_G) = g$, we must have $(L_g\phi)^{-1}(g) = 1_G$. Now observe the following:

$$(L_g\phi)^{-1}\beta L_g\phi(1_G) = (L_g\phi)^{-1}(\beta(g)) = (L_g\phi)^{-1}(g) = 1_G$$

Therefore, $(L_g\phi)^{-1}\beta L_g\phi$ fixes the identity vertex 1_G in Γ. Since $L_g\phi(\phi^{-1}(t)) = gt$, we must have $(L_g\phi)^{-1}(gt) = \phi^{-1}(t)$. Now observe the following:

$$(L_g\phi)^{-1}\beta L_g\phi(\phi^{-1}(t)) = (L_g\phi)^{-1}(\beta(gt)) = (L_g\phi)^{-1}(gt) = \phi^{-1}(t)$$

Therefore, $(L_g\phi)^{-1}\beta L_g\phi$ fixes the vertex $\phi^{-1}(t)$ in Γ. Since $\phi \in \mathrm{Aut}(G,\mathcal{S}^{\pm})$, $\phi^{-1}(t) \in \mathcal{S}^{\pm}$, and so $e(1_G,\phi^{-1}(t))$ is an edge in Γ. The preceding work shows that $(L_g\phi)^{-1}\beta L_g\phi$ fixes the edge $e(1_G,\phi^{-1}(t))$. However, the only element of G_R which fixes an edge incident to 1_G in Γ is id_Γ, as we've already shown. So $(L_g\phi)^{-1}\beta L_g\phi = \mathrm{id}_\Gamma$, implying that $\beta L_g\phi = L_g\phi$, which further implies that $\beta = \mathrm{id}_\Gamma$. Therefore, if $\beta \in G_R$ is arbitrary and if β fixes any edge in Γ, then

$\beta = \mathrm{id}_\Gamma$.

We can now conclude that the action of G_R on Γ is edge free. So G_R is a reflection group and $G_R \subseteq \mathrm{Aut}_{\mathrm{c.p.}}(\Gamma)$. Moreover, for any $s \in \mathcal{S}^{\pm}$, the reflection $L_s\phi_s \in G_R$ inverts the edge $e(1_G, s)$, by assumption. By applying Proposition 7.4, we can conclude that $(G, \mathcal{S})$ is strongly totally reflected.

□

If $(G, \mathcal{S})$ is totally reflected, then we know that for each $e \in E(\Gamma)$ we can find a color-preserving reflection acting on $\Gamma = \Gamma(G, \mathcal{S})$ which inverts the edge e. In particular, for each $s \in \mathcal{S}^{\pm}$ there exists a color-preserving reflection on Γ which inverts the edge $e(1_G, s)$. Fix $t \in \mathcal{S}$. We know that the color-preserving reflection on Γ which inverts the edge $e(1_G, t)$ must specifically have the form $L_t\phi_t$, where $\phi_t \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$ satisfies $\phi_t(t) = t^{-1}$ and $(\phi_t)^2 = \mathrm{id}_G$. Likewise, the color-preserving reflection on Γ which inverts the edge $e(1_G, t^{-1})$ must specifically have the form $L_{t^{-1}}\phi_{t^{-1}}$, where $\phi_{t^{-1}} \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$ satisfies $\phi_{t^{-1}}(t^{-1}) = t$ and $(\phi_{t^{-1}})^2 = \mathrm{id}_G$. If we conjugate $L_t\phi_t$ by the graph automorphism $L_{t^{-1}}$, we obtain another color-preserving reflection on Γ which inverts the edge $L_{t^{-1}}\big(e(1_G, t)\big) = e(t^{-1}, t)$. Therefore, this conjugate reflection also inverts $\overline{e(t^{-1}, t)} = e(1_G, t^{-1})$. Observe the following:

$$\begin{aligned}(L_{t^{-1}})L_t\phi_t(L_{t^{-1}})^{-1}(x) &= L_{t^{-1}}L_t\phi_t L_t(x) = \phi_t L_t(x) = \phi_t(tx)\\ &= \phi_t(t)\phi_t(x) = t^{-1}x = L_{t^{-1}}\phi_t(x)\end{aligned}$$

Therefore, $L_{t^{-1}}\phi_t$ is a color-preserving reflection on Γ, possibly different than $L_{t^{-1}}\phi_{t^{-1}}$, which inverts the edge $e(1_G, t^{-1})$. It may be convenient, and even necessary, at times to assume that $\phi_{t^{-1}} = \phi_t$. In fact, we will make such an assumption in our next proposition, and the proof that we give for the proposition will depend on this added condition.

Another condition that must be satisfied in order to state and prove the next proposition involves the ϕ_s component of the color-preserving reflection $L_s\phi_s$ on $\Gamma = \Gamma(G, \mathcal{S})$. More generally, if $L_g\phi$ is any color-preserving automorphism on a Cayley graph $\Gamma(G, \mathcal{S})$, where $n := |\mathcal{S}|$, then we will say that ϕ satisfies the **$(\mathbb{Z}_2)^n$ Condition** if for any $s \in \mathcal{S}$ we have $\phi(s) = s$ or $\phi(s) = s^{-1}$. Note that to be a color-preserving automorphism, ϕ must already be an element of $\mathrm{Aut}(G, \mathcal{S}^{\pm})$. However, an arbitrary member of $\mathrm{Aut}(G, \mathcal{S}^{\pm})$ may not satisfy the $(\mathbb{Z}_2)^n$ Condition. Let H be the set of those elements of $\mathrm{Aut}(G)$ which satisfy the $(\mathbb{Z}_2)^n$ Condition. One can check that H is a group. If $\mathcal{S} = \{s_1, ..., s_n\}$, then for each $i \in I_n$ define an element α_i of H by stating that $\alpha_i(s_j) = s_j$ if $j \neq i$ and $\alpha_i(s_i) = s_i^{-1}$. Then the set $\{\alpha_i \mid i \in I_n\}$ generates H. We can define a function $\Phi : H \to (\mathbb{Z}_2)^n$ by sending the generator α_i of H to the element of $(\mathbb{Z}_2)^n$ which has a 1 in the i^{th} position and a 0 in every other position. One can check that Φ is an isomorphism. Therefore, an element $\phi \in \mathrm{Aut}(G, \mathcal{S}^{\pm})$ satisfies the $(\mathbb{Z}_2)^n$ Condition if and only if $\phi \in H \cong (\mathbb{Z}_2)^n$.

The next proposition will prove to be a useful tool in showing that certain group systems are strongly totally reflected.

Proposition 7.6. *Let $(G, \mathcal{S})$ be a totally reflected group system. Let $n := |\mathcal{S}|$. Label the n distinct elements of $\mathcal{S}$ as $s_1, ..., s_n$. For each $s \in \mathcal{S}^{\pm}$, let $L_s\phi_s$ be a color-preserving reflection on $\Gamma = \Gamma(G, \mathcal{S})$ inverting the edge $e(1_G, s)$. Assume that $\phi_s = \phi_{s^{-1}}$ for all $s \in \mathcal{S}$. If all relations for the system $(G, \mathcal{S})$ have even exponent sum for each $s \in \mathcal{S}$ and if each ϕ_s satisfies the $(\mathbb{Z}_2)^n$ Condition, then $(\star\star)$ will hold, and $(G, \mathcal{S})$ will be strongly totally reflected.*

Proof. Let $s_{i_1}s_{i_2}\cdots s_{i_k}$ be an arbitrary (nonempty) word in $\mathcal{S}^{\pm}$ such that the antecedent of condition $(\star\star)$ holds, that is,

$$s_{i_1}\cdot[\phi_{s_{i_1}}(s_{i_2})]\cdot[\phi_{s_{i_1}}\phi_{s_{i_2}}(s_{i_3})]\cdots[\phi_{s_{i_1}}\cdots\phi_{s_{i_{k-1}}}(s_{i_k})]=1_G$$

Assume that for each $s\in\mathcal{S}$, ϕ_s satisfies the $(\mathbb{Z}_2)^n$ Condition. Then the preceding equation can be rewritten as shown below, where each $\epsilon_i\in\{-1,1\}$:

$$(s_{i_1})^{\epsilon_1}(s_{i_2})^{\epsilon_2}\cdots(s_{i_k})^{\epsilon_k}=1_G$$

Assume, also, that all relations for the system $(G,\mathcal{S})$ have even exponent sum for each $s\in\mathcal{S}$. Let p_s be the integer obtained by summing those values of ϵ_j for which $s_{i_j}=s$ or s^{-1}, where $s\in\mathcal{S}$. In other words, p_s is the exponent sum for the generator s in the relation given in the preceding equation. Then by assumption, each p_s is an even integer (possibly zero). For each $s\in\mathcal{S}$, let q_s be the integer obtained by summing those values of $|\epsilon_j|$ for which $s_{i_j}=s$ or s^{-1}. In other words, q_s counts how many times s or s^{-1} occurs in the product $(s_{i_1})^{\epsilon_1}(s_{i_2})^{\epsilon_2}\cdots(s_{i_k})^{\epsilon_k}$. Note that q_s must be a (nonnegative) even integer, since p_s is even.

Since ϕ_s satisfies the $(\mathbb{Z}_2)^n$ Condition for each $s\in\mathcal{S}$, the elements of the set $\{\phi_s \mid s\in\mathcal{S}\}$ commute with one another. Therefore, we can observe the following:

$$\phi_{s_{i_1}}\cdots\phi_{s_{i_k}}=(\phi_{s_1})^{q_{s_1}}(\phi_{s_2})^{q_{s_2}}\cdots(\phi_{s_n})^{q_{s_n}}$$

However, each ϕ_{s_i} must have order 2, since $L_{s_i}\phi_{s_i}$ is a reflection. Therefore, we can now conclude that $(\phi_{s_1})^{q_{s_1}}(\phi_{s_2})^{q_{s_2}}\cdots(\phi_{s_n})^{q_{s_n}}=\mathrm{id}_G$, since for each $i\in I_n$, $(\phi_{s_i})^{q_{s_i}}=(\phi_{s_i})^{q_{s_i}(\bmod 2)}=(\phi_{s_i})^0=\mathrm{id}_G$.

Therefore, Condition $(\star\star)$ is satisfied, implying that $(G, \mathcal{S})$ is strongly totally reflected.

□

Example 7.7. Suppose that G is a finitely generated abelian group with invariant-factor decomposition given by $\mathbb{Z}_{n_1} \times \mathbb{Z}_{n_2} \times \cdots \times \mathbb{Z}_{n_k} \times \mathbb{Z}^r$, where $n_1, ..., n_k$, k, and r are nonnegative integers with $n_i \geq 2$ for all $i \in I_k$ and with $n_i | n_{i+1}$ for $i \in I_{k-1}$ in the event that $k \geq 2$. For simplicity, let us consider the case where $k, r > 0$. If we have any hope of G being s.t.r., then G must be t.r. In Chapter 5 we showed that G is t.r. if and only if n_1 is even. Therefore, assume n_1 is even.

For each $i \in I_k$, $\mathbb{Z}_{n_i}$ is a finite cyclic group and thus can be generated by a single element of order n_i. Let a_i denote this element. Then $\mathbb{Z}_{n_i} = \langle a_i \mid (a_i)^{n_i} = 1_{\mathbb{Z}_{n_i}} \rangle$. Now $\mathbb{Z}$ is an infinite cyclic group and thus can be generated by a single torsion-free element. However, for each of the r factors of $\mathbb{Z}$ that appears in the product $\mathbb{Z}^r$, we must use a different generator. Therefore, we will say that the jth factor of $\mathbb{Z}$ in $\mathbb{Z}^r$ is generated by the element b_j. That is, $\mathbb{Z}^r = \langle b_1, ..., b_r \mid [b_i, b_j] = 1_{\mathbb{Z}^r} \text{ for all } i, j \in I_r, i \neq j \rangle$. We can now write a presentation for the finitely generated abelian group G as shown here:

$$G = \langle a_1, ..., a_k, b_1, ..., b_r \mid (a_i)^{n_i} = [a_i, a_j] = [b_p, b_q] = [a_i, b_p] = 1_G,$$
$$\text{for all } i, j \in I_k, i \neq j; p, q \in I_r, p \neq q \rangle$$

Let $\mathcal{S} := \{a_1, ..., a_k, b_1, ..., b_r\}$

Any relation in a group of the form $[x, y] = xyx^{-1}y^{-1} = 1$, where x and y are generators for the group, has exponent sum of zero for x and for y.

It also has exponent sum of zero for every other generator in the group, trivially. For any element z of the group, the relation $zz^{-1} = 1$ will have exponent sum of zero for all of the group generators. Therefore, if a group has a presentation in which all the defining relations are commutators, then *any* relation in the group will have even (specifically, zero) exponent sum for each group generator. More generally, if a group has a presentation in which all the defining relations have even exponent sum for each generator, then an arbitrary relation in the group will have even exponent sum for each group generator.

In our group G with presentation given above, we have defining relations which are commutators and defining relations of the form $(a_i)^{n_i} = 1_G$ for each $i \in I_k$. Every n_i is even since n_1 is even and since $n_i | n_{i+1}$ for all $i \in I_{k-1}$ in the event that $k \geq 2$. Therefore, all of our defining relations, and consequently *all* relations in the group G, have even exponent sum for each $s \in \mathcal{S}$.

Fix $i \in I_k$. We know from previous work that $(\mathbb{Z}_{n_i}, \{a_i\})$ is totally reflected because n_i is even. Let ρ_{a_i} be a group automorphism on $\mathbb{Z}_{n_i}$ defined by stating that $\rho(a_i) = a_i^{-1}$. Then we know from previous work that $L_{a_i}\rho_{a_i}$ and $L_{a_i^{-1}}\rho_{a_i}$ are color-preserving reflections on $\Gamma(\mathbb{Z}_{n_i}, \{a_i\})$ which invert the edges $e(1_{\mathbb{Z}_{n_i}}, a_i)$ and $e(1_{\mathbb{Z}_{n_i}}, a_i^{-1})$, respectively. Now define a group automorphism ϕ_{a_i} on G as shown here:

$$\text{for } s \in \mathcal{S},\ \phi_{a_i}(s) \coloneqq \begin{cases} s^{-1}, \text{ if } s = a_i \\ \\ s, \text{ if } s \neq a_i \end{cases}$$

We know from our previous work involving direct products of t.r. groups that $L_{a_i}\phi_{a_i}$ and $L_{a_i^{-1}}\phi_{a_i}$ are color-preserving reflections on $\Gamma = \Gamma(G, \mathcal{S})$ which invert the edges $e(1_G, a_i)$ and $e(1_G, a_i^{-1})$, respectively.

Now fix $p \in I_r$. We know from previous work that $(\mathbb{Z}, \{b_p\})$ is totally reflected. Let ρ_{b_p} be a group automorphism on $\mathbb{Z}$ defined by stating the $\rho(b_p) = b_p^{-1}$. Then $L_{b_p}\rho_{b_p}$ and $L_{b_p^{-1}}\rho_{b_p}$ are color-preserving reflections on $\Gamma(\mathbb{Z}, \{b_p\})$ which invert the edges $e(1_{\mathbb{Z}}, b_p)$ and $e(1_{\mathbb{Z}}, b_p^{-1})$, respectively. If we extend ρ_{b_p} to a group automorphism ϕ_{b_p} on G analogously to the way we extended ρ_{a_i} to ϕ_{a_i}, then we know that $L_{b_p}\phi_{b_p}$ and $L_{b_p^{-1}}\phi_{b_p}$ are color-preserving reflections on $\Gamma = \Gamma(G, \mathcal{S})$ which invert the edges $e(1_G, b_p)$ and $e(1_G, b_p^{-1})$, respectively.

We know that $(G, \mathcal{S})$ is totally reflected. For any $s \in \mathcal{S}^{\pm}$, $L_s\phi_s$ is a color-preserving reflection on $\Gamma = \Gamma(G, \mathcal{S})$ which inverts the edge $e(1_G, s)$. Because of the way we defined ρ_s and then ϕ_s for each $s \in \mathcal{S}$, it is clear that ϕ_s satisfies the $(\mathbb{Z}_2)^n$ condition for each $s \in \mathcal{S}$. As we already noted, every relation in the group G has even exponent for each $s \in \mathcal{S}$. Therefore, by Proposition 7.6, we can conclude that $(G, \mathcal{S})$ is strongly totally reflected.

Similar arguments can be used to show that $(G, \mathcal{S})$ is s.t.r. even when $k = 0$ or $r = 0$ in the invariant factor decomposition for G, as long as $n_1 \geq 2$ is even in the event that $k > 0$. Therefore, *any* finitely generated abelian group with even first invariant factor is strongly totally reflected.

◇

7.3 The Case of Right-Angled Artin Groups

Suppose Δ is a simplicial graph on n vertices $\{v_i \mid i \in I_n\} = V(\Delta) =: \mathcal{S}$. Let A_Δ denote the right-angled Artin group determined by Δ. A standard way of defining A_Δ is in terms of a presentation such as the one shown here:

$$A_\Delta := \langle v_1, \dots, v_n \,\big|\, [v_i, v_j] = 1 \iff \{v_i, v_j\} \in E(\Delta)\rangle$$

where, recall, $\{v_i, v_j\}$ represents the edge in Δ between v_i and v_j. Consider the collection of groups $\mathcal{G} = \{G_i \,\big|\, i \in I_n\}$, where for each $i \in I_n$ we have $G_i = \langle \mathcal{S}_i \,\big|\, \mathcal{R}_i\rangle$ with $\mathcal{S}_i = \{v_i\}$ and $\mathcal{R}_i = \varnothing$. In other words, $G_i = \langle v_i \rangle$ is the infinite cyclic group. Then we can easily observe that $A_\Delta = \overline{G}(\Delta, \mathcal{G})$, where $\overline{G}(\Delta, \mathcal{G})$ is the presentation-form of the right-angled product as defined in Chapter 6.

Fix an arbitrary $i \in I_n$. By Example 5.4, we know that $(G_i, \mathcal{S}_i)$ is a totally reflected system. Therefore, there exists a color-preserving graph reflection on $\Gamma(G_i, \mathcal{S}_i)$ which inverts the edge $e(1, v_i)$. By Example 4.11, we know that this graph reflection must specifically be given by $L_{v_i}\rho_{v_i}$, where $\rho_{v_i} : G_i \to G_i$ is the automorphism defined by stating that $\rho_{v_i}(v_i) = v_i^{-1}$.

Using the process from the proof of Lemma 6.8, we can extend ρ_{v_i} to a function $\phi_{v_i} : A_\Delta \to A_\Delta$ as defined here:

$$\phi_{v_i}(v) := \begin{cases} v^{-1}, & \text{if } \ v \in \{v_i, v_i^{-1}\} \\ v, & \text{if } \ v \in \mathcal{S}^\pm - \{v_i, v_i^{-1}\} \end{cases}$$

Our work in the proof of Lemma 6.8 gives us that $\phi_{v_i} : A_\Delta \to A_\Delta$ is an automorphism and $\phi_{v_i} \in \mathrm{Aut}(A_\Delta, \mathcal{S}^\pm)$. Moreover, $(\phi_{v_i})^2 = \mathrm{id}_{A_\Delta}$, since $L_{v_i}\rho_{v_i}$ being a reflection on $\Gamma(G_i, \mathcal{S}_i)$ implies that $(\rho_{v_i})^2 = \mathrm{id}_{G_i}$. Lemma 6.8 also tells us that $L_{v_i}\phi_{v_i}$ is a reflection on $\Gamma(A_\Delta, \mathcal{S})$ which inverts the edge $e(1, v_i)$. Using an analogous argument to the one just given, we also can conclude that $L_{v_i^{-1}}\phi_{v_i}$ is a reflection on $\Gamma(A_\Delta, \mathcal{S})$ which inverts the edge $e(1, v_i^{-1})$ in $\Gamma(A_\Delta, \mathcal{S})$.

Since the defining relations of A_Δ are commutators, all relations for the system $(A_\Delta, \mathcal{S})$ will have an even exponent sum for each $s \in \mathcal{S}$. Moreover, it is clear from its definition that ϕ_{v_i} satisfies the $(\mathbb{Z}_2)^n$ Condition for each $i \in I_n$. Therefore, by Proposition 7.6, $(A_\Delta, \mathcal{S})$ is a strongly totally reflected system.

Let $R := \{L_{v_i}\phi_{v_i} \mid i \in I_n\} \cup \{L_{v_i^{-1}}\phi_{v_i} \mid i \in I_n\}$. From the preceding discussion, we can see that every element of R is a reflection on $\Gamma = \Gamma(A_\Delta, \mathcal{S})$. Let A_r denote the group generated by R, that is, $A_r := \langle R \rangle$. Then A_r is a reflection group acting on Γ; $A_r \subseteq \mathrm{Aut}_{\text{c.p.}}(\Gamma)$ since each generator of A_r is an element of $\mathrm{Aut}_{\text{c.p.}}(\Gamma)$; for any $v_i \in \mathcal{S}$ the reflection $L_{v_i}\phi_{v_i}$ inverts the edge $e(1_{A_\Delta}, v_i)$; and for any $v_i \in \mathcal{S}$ the reflection $L_{v_i^{-1}}\phi_{v_i}$ inverts the edge $e(1_{A_\Delta}, v_i^{-1})$. Proposition 7.4 allows us to conclude that for *any* edge e in Γ there exists a reflection in A_r which inverts the edge e. The next proposition establishes a strong relationship between the group A_r and the original right-angled Artin group A_Δ.

Proposition 7.8. *The groups A_Δ and A_r are commensurable.*

Proof. In order to show that A_Δ and A_r are commensurable, we must show that $A_\Delta \cap A_r$ is a finite-index subgroup in both A_Δ and A_r. We will actually prove that there exists a group G containing both A_Δ and A_r as subgroups such that $[G : A_\Delta] < \infty$ and $[G : A_r] < \infty$. Then basic group theory will allow us to conclude the following:

$$[A_\Delta : A_\Delta \cap A_r] \le [G : A_r] < \infty \quad \text{and} \quad [A_r : A_\Delta \cap A_r] \le [G : A_\Delta] < \infty$$

Consider the group generated by the set $\{\phi_{v_i} \mid i \in I_n\}$, where each ϕ_{v_i} is as defined earlier in this section. We can easily observe that this group is isomorphic to $(\mathbb{Z}_2)^n$. In this proof, we will think of $(\mathbb{Z}_2)^n$ as the group $\left\langle\{\phi_{v_i} \mid i \in I_n\}\right\rangle$. Recall that A_Δ can be thought of as a subgroup of $\mathrm{Aut}(A_\Delta)$, as established by Lemma 3.5 and Corollary 3.6, where each element $a \in A_\Delta$ is identified with L_a.

Let $G := A_\Delta \rtimes (\mathbb{Z}_2)^n$. Elements of G are of the form $L_a\alpha$, where $a \in A_\Delta$ and $\alpha \in (\mathbb{Z}_2)^n$. Note that A_Δ and $(\mathbb{Z}_2)^n$ are both subgroups of G. Consider a function $\theta : A_\Delta \rtimes (\mathbb{Z}_2)^n \to (\mathbb{Z}_2)^n$ defined by stating that for $a \in A_\Delta$ and $\alpha \in (\mathbb{Z}_2)^n$, $\theta(L_a\alpha) := \alpha$. Suppose $a_1, a_2 \in A_\Delta$ and $\alpha_1, \alpha_2 \in (\mathbb{Z}_2)^n$ are arbitrary. Then observe the following:

$$\begin{aligned}
\theta(L_{a_1}\alpha_1 L_{a_2}\alpha_2) &= \theta\left(L_{a_1\alpha_1(a_2)}\alpha_1\alpha_2\right) \\
&= \alpha_1\alpha_2 \\
&= \theta(L_{a_1}\alpha_1)\,\theta(L_{a_2}\alpha_2)
\end{aligned}$$

Therefore, θ is a homomorphism. We can see that $\ker\theta = A_\Delta$ and $\mathrm{im}\,\theta = (\mathbb{Z}_2)^n$. Then basic group theory tells us that $[G : \ker\theta] = |\mathrm{im}\,\theta|$, which implies that $[G : A_\Delta] = 2^n < \infty$.

Suppose $L_a\alpha \in A_\Delta \rtimes (\mathbb{Z}_2)^n = G$. Since $a \in A_\Delta$, a can be written as a word in $\mathcal{S}^\pm$, say $a = u_1 \cdots u_m$, where $u_i \in \mathcal{S}^\pm$ for each $i \in I_m$. Let $\epsilon_1 = 1$ and for integers $i \geq 2$ define ϵ_i as shown here:

$$\epsilon_i := \begin{cases} 1, & \text{if } \exists \text{ an even number of indices } j \in I_{i-1} \text{ such that } u_j = u_i \text{ or } u_i^{-1} \\ -1, & \text{if } \exists \text{ an odd number of indices } j \in I_{i-1} \text{ such that } u_j = u_i \text{ or } u_i^{-1} \end{cases}$$

For each $e \in I_m$ let $w_i := u_i^{\epsilon_i}$. Therefore, for each $i \in I_m$ with $i \geq 2$, $\phi_{w_1}\phi_{w_2} \cdot \cdot \cdot \phi_{w_{i-1}}(w_i) = u_i$. Also, $w_1 = u_1$, since $\epsilon_1 = 1$. Note that $L_{w_i}\phi_{w_i} \in A_r$ for each $i \in I_m$. Then $L_{w_1}\phi_{w_1}L_{w_2}\phi_{w_2} \cdots L_{w_m}\phi_{w_m} \in A_r$, and for any $x \in A_\Delta$ we can see the following:

$$
\begin{aligned}
L_{w_1}\phi_{w_1} \cdots L_{w_m}\phi_{w_m}(x) &= w_1\phi_{w_1}(w_2) \cdots \phi_{w_1}\phi_{w_2}\cdots\phi_{w_{m-1}}(w_m)\phi_{w_1}\cdots\phi_{w_m}(x) \\
&= u_1u_2\cdots u_m\phi_{w_1}\cdots\phi_{w_m}(x) \\
&= a\phi_{w_1}\cdots\phi_{w_m}(x) \\
&= L_a\phi_{w_1}\cdots\phi_{w_m}(x)
\end{aligned}
$$

We can see then that $L_a\alpha = (L_a\phi_{w_1} \cdots \phi_{w_m})(\phi_{w_m} \cdots \phi_{w_1}\alpha)$, where $L_a\phi_{w_1} \cdot \cdot \cdot \phi_{w_m} = L_{w_1}\phi_{w_1} \cdots L_{w_m}\phi_{w_m} \in A_r$, as proven, and $\phi_{w_m} \cdots \phi_{w_1}\alpha \in (\mathbb{Z}_2)^n$. This implies that $L_a\alpha \in A_r\beta$ for some $\beta \in (\mathbb{Z}_2)^n$. Consequently, there are at most $|(\mathbb{Z}_2)^n| = 2^n$ cosets of A_r in G, that is, $[G : A_r] \leq 2^n < \infty$.

Since $[G : A_r] < \infty$ and $[G : A_\Delta] < \infty$, we now can conclude that $[A_\Delta : A_\Delta \cap A_r] < \infty$ and $[A_r : A_\Delta \cap A_r] < \infty$, for reasons explained previously.

Therefore, A_Δ and A_r are commensurable.

□

The chambers of A_r acting on $\Gamma = \Gamma(A_\Delta, \mathcal{S})$ are the vertices of Γ since for any edge there exists a reflection in A_r which inverts that edge. Fix the vertex 1_{A_Δ} to be the base vertex. Then the set of fundamental reflections is what we are calling R. We know now, as a result of our discussion earlier in this section, that (A_r, R) is a reflection system on the graph Γ. Then by Theorem 4.16 we can conclude that (A_r, R) forms a Coxeter system where for any $r_i, r_j \in R$, m_{r_i,r_j} is taken to be the order of r_ir_j in A_r.

Recall that $R = \{L_{v_i}\phi_{v_i} \mid i \in I_n\} \cup \{L_{v_i^{-1}}\phi_{v_i} \mid i \in I_n\}$. For convenience, for each $i \in I_n$, let $r_i = L_{v_i}\phi_{v_i}$ and let $t_i = L_{v_i^{-1}}\phi_{v_i}$. Since the elements of R are reflections, we know that $m_{r_i,r_i} = m_{t_i,t_i} = 1$ for any $i \in I_n$. Assume that $\{v_i, v_j\} \in E(\Delta)$. Then $v_i \neq v_j$. Moreover, $[v_i, v_j] = v_i v_j v_i^{-1} v_j^{-1} = 1_{A_\Delta}$, implying that v_i and v_j commute. From this we can observe that v_i and v_j^{-1} commute, v_i^{-1} and v_j commute, and v_i^{-1} and v_j^{-1} commute, also. Suppose $x \in A_\Delta = V(\Gamma)$ is arbitrary. Then observe the following:

$$
\begin{aligned}
r_i r_j(x) &= L_{v_i}\phi_{v_i} L_{v_j}\phi_{v_j}(x) = v_i \phi_{v_i}(v_j)\phi_{v_i}\phi_{v_j}(x) = v_i v_j \phi_{v_i}\phi_{v_j}(x) \\
&= v_j v_i \phi_{v_j}\phi_{v_i}(x) = v_j \phi_{v_j}(v_i)\phi_{v_j}\phi_{v_i}(x) = L_{v_j}\phi_{v_j} L_{v_i}\phi_{v_i}(x) \quad = r_j r_i(x)
\end{aligned}
$$

Therefore, from the above argument and from the fact that every element of R has order 2, we can conclude that $(r_i r_j)^2 = \mathrm{id}_\Gamma$. Similar work would allow us to show that $(t_i t_j)^2 = \mathrm{id}_\Gamma$ and $(r_i t_j)^2 = \mathrm{id}_\Gamma$ as well. Therefore, for any $i, j \in I_n$, $i \neq j$, such that $\{v_i, v_j\} \in E(\Delta)$, $m_{r_i,r_j} = m_{t_i,t_j} = m_{r_i,t_j} = 2$.

Assume now that $i, j \in I_n$, $i \neq j$, such that $\{v_i, v_j\} \notin E(\Delta)$. Then in A_Δ, no relation exists between v_i and v_j. In particular, $[v_i, v_j] \neq 1_{A_\Delta}$. We claim that for any $n \in \mathbb{N}$ and for $x \in A_\Delta$ arbitrary, the following statements are true:

$$
\begin{aligned}
(r_i r_j)^{2n-1}(x) &= L_{[v_i,v_j]^{n-1} v_i v_j}\phi_{v_i}\phi_{v_j}(x) \\
(r_i r_j)^{2n}(x) &= L_{[v_i,v_j]^n}(x)
\end{aligned}
$$

We will prove this claim by induction. First, observe the following:

$$
r_i r_j(x) = L_{v_i}\phi_{v_i} L_{v_j}\phi_{v_j}(x) = v_i \phi_{v_i}(v_j)\phi_{v_i}\phi_{v_j}(x)
$$

$$= v_i v_j \phi_{v_i}\phi_{v_j}(x) = L_{v_i v_j}\phi_{v_i}\phi_{v_j}(x)$$

Additionally,

$$\begin{aligned}(r_i r_j)^2(x) &= r_i r_j\left(r_i r_j(x)\right) = L_{v_i v_j}\phi_{v_i}\phi_{v_j}\left(L_{v_i v_j}\phi_{v_i}\phi_{v_j}(x)\right)\\ &= v_i v_j v_i^{-1} v_j^{-1}\phi_{v_i}\phi_{v_j}\phi_{v_i}\phi_{v_j}(x) = L_{[v_i,v_j]}\left(\phi_{v_i}\right)^2\left(\phi_{v_j}\right)^2(x)\\ &= L_{[v_i,v_j]}(x)\end{aligned}$$

This shows that the statements are true when $n = 1$. Suppose that the statements are true for some $k \in \mathbb{N}$. Then observe the following:

$$\begin{aligned}(r_i r_j)^{2(k+1)-1}(x) &= (r_i r_j)^{2k-1}(r_i r_j)^2(x) = (r_i r_j)^{2k-1}\left(L_{[v_i,v_j]}(x)\right)\\ &= L_{[v_i,v_j]^{k-1}v_i v_j}\phi_{v_i}\phi_{v_j}\left(L_{[v_i,v_j]}(x)\right)\\ &= [v_i, v_j]^{k-1} v_i v_j \phi_{v_i}\phi_{v_j}(v_i v_j v_i^{-1} v_j^{-1})\phi_{v_i}\phi_{v_j}(x)\\ &= [v_i, v_j]^{k-1} v_i v_j v_i^{-1} v_j^{-1} v_i v_j \phi_{v_i}\phi_{v_j}(x)\\ &= L_{[v_i,v_j]^k v_i v_j}\phi_{v_i}\phi_{v_j}(x)\\ (r_i r_j)^{2(k+1)}(x) &= (r_i r_j)^{2k}(r_i r_j)^2(x) = (r_i r_j)^{2k}\left(L_{[v_i,v_j]}(x)\right)\\ &= L_{[v_i,v_j]^k} L_{[v_i,v_j]}(x)\\ &= L_{[v_i,v_j]^{k+1}}(x)\end{aligned}$$

This shows that by using the inductive hypothesis for $n = k$, along with the initial step, we can show that the statements are true for $n = k + 1$. Therefore, the statements are true for any $n \in \mathbb{N}$, where $x \in A_\Delta$ is arbitrary. In particular, the statements must be true for $x = 1_{A_\Delta}$. That is, for any $n \in \mathbb{N}$, $(r_i r_j)^{2n-1}(1_{A_\Delta}) = [v_i, v_j]^{n-1} v_i v_j$, since $\phi_{v_i}(1_{A_\Delta}) = \phi_{v_j}(1_{A_\Delta}) = 1_{A_\Delta}$, and

$(r_i r_j)^{2n}(1_{A_\Delta}) = [v_i, v_j]^n$.

Note that for no $n \in \mathbb{N}$ can $(r_i r_j)^{2n-1} = \mathrm{id}_\Gamma$ or $(r_i r_j)^{2n} = \mathrm{id}_\Gamma$. If there was an $n \in \mathbb{N}$ such that $(r_i r_j)^{2n-1}$ or $(r_i r_j)^{2n}$ were equal to id_Γ, then we would have that $[v_i, v_j]^{n-1} v_i v_j = 1_{A_\Delta}$ or $[v_i, v_j]^n = 1_{A_\Delta}$. However, neither of these scenarios can happen since $\{v_i, v_j\} \notin E(\Delta)$. Therefore, for any $m \in \mathbb{N}$, $(r_i r_j)^m \neq \mathrm{id}_\Gamma$. Similar work would allow us to show that under the same circumstances involving i and j, for any $m \in \mathbb{N}$ we have $(r_i t_j)^m \neq \mathrm{id}_\Gamma$ and $(t_i t_j)^m \neq \mathrm{id}_\Gamma$. Therefore, for any $i, j \in I_n$, $i \neq j$, such that $\{v_i, v_j\} \notin E(\Delta)$, $m_{r_i, r_j} = m_{t_i, t_j} = m_{r_i, t_j} = \infty$.

We now can write the Coxeter presentation for A_r:

$$A_r = \langle r_1, \ldots, r_n, t_1, \ldots, t_n \mid r_i^2 = t_i^2 = 1_{A_r} = \mathrm{id}_\Gamma \ \forall i \in I_n;\ \text{for any } i, j \in I_n, i \neq j,$$
$$(r_i r_j)^2 = (r_i t_j)^2 = (t_i t_j)^2 = 1_{A_r} = \mathrm{id}_\Gamma \iff \{v_i, v_j\} \in E(\Delta) \rangle$$

We can see from this presentation that A_r is not only a Coxeter group but a *right-angled* Coxeter group. Proposition 7.8 can then be interpreted as saying that any right-angled Artin group is commensurable with a right-angled Coxeter group. Though this result is already known [3], our approach to proving it using the theory of totally reflected groups is novel.

Example 7.9. Suppose that A_Δ is a right-angled Artin group determined by the graph Δ shown in Figure 7.2.

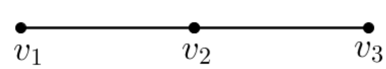

Figure 7.2: A graph Δ determining a right-angled Artin group A_Δ

We can view A_Δ as being the right-angled product determined by Δ

and $\mathcal{G}$, where $\mathcal{G} = \{G_1, G_2, G_3\}$ and $G_i = \langle v_i \rangle$ is an infinite cyclic group for each $i \in I_3$. The associated reflection group A_r, defined in the manner discussed prior to Proposition 7.8, is a right-angled Coxeter group. Since any Coxeter group can be described as a right-angled product, there exists a graph Δ' and a collection of groups $\mathcal{G}'$ such that A_r is the right-angled product determined by Δ' and $\mathcal{G}'$. For each $i \in I_3$, let $r_i := L_{v_i}\phi_{v_i}$, $t_i := L_{v_i^{-1}}\phi_{v_i}$, $H_i := \langle r_i \mid (r_i)^2 = 1_{H_i} \rangle$, and $K_i := \langle t_i \mid (t_i)^2 = 1_{K_i} \rangle$. Take $\mathcal{G}'$ to be the collection of groups $\{H_1, H_2, H_3, K_1, K_2, K_3\}$. The graph Δ' is depicted in Figure 7.3.

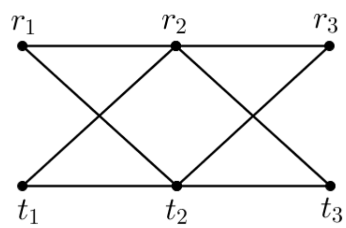

Figure 7.3: The graph Δ' determining the right-angled Coxeter group A_r

The graph Δ that we used here to define the right-angled Artin group A_Δ is quite basic. However, treating this example as a guide, we can easily imagine how to obtain the graph Δ' from the graph Δ in the general setting.

◇

Bibliography

[1] H. S. M. Coxeter. Discrete groups generated by reflections. *Ann. of Math. (2)*, 35(3):588–621, 1934.

[2] H. S. M. Coxeter. The complete enumeration of finite groups of the form (ri) ˆ (2)=(rirj) ˆ (kij)=1. *Journal of the London Mathematical Society*, s1-10(1):21–25, 1935.

[3] Michael W. Davis and Tadeusz Januszkiewicz. Right-angled Artin groups are commensurable with right-angled Coxeter groups. *J. Pure Appl. Algebra*, 153(3):229–235, 2000.

[4] James E. Humphreys. *Reflection groups and Coxeter groups*, volume 29 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1990.

[5] Thomas W. Hungerford. *Algebra*, volume 73 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1980. Reprint of the 1974 original.

[6] Wilhelm Magnus, Abraham Karrass, and Donald Solitar. *Combinatorial group theory*. Dover Publications, Inc., New York, revised edition, 1976.

[7] William S. Massey. *A basic course in algebraic topology*, volume 127 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991.

[8] A. Miller. Coxeter groups and the Davis construction. University of Oklahoma, Department of Mathematics, 1994. Course lecture notes.

[9] A. Miller. Graph reflection groups. University of Oklahoma, Department of Mathematics, 2007. Course lecture notes.

[10] Ric Wade. Combinatorial group theory notes, 2014. Available online at http://www.math.utah.edu/~wade/5520spring2014/Groups.pdf.