On Groups and Their Graphs

by

Melissa U. Sherman-Bennett

A Thesis submitted to the Faculty

in partial fulfillment of the requirements for the

BACHELOR OF ARTS

Accepted

William Dunbar, Thesis Advisor

Aaron Williams, Second Reader

Patrick Dragon, Third Reader

Ian Bickford, Provost and Vice President

Bard College at Simon's Rock
Great Barrington, MA

May, 2016

Notation

Symbol	Meaning
$A \subseteq B$	A is a subset of B
$A \subsetneq B$	A is a proper subset of B
$X \subseteq Y$	X is a subgraph of Y
$H \leq G$	H is a subgroup of G
$H \lneq G$	H is a proper subgroup of G
$H \trianglelefteq G$	H is a normal subgroup of G
$H \cong G$	the groups H and G are isomorphic
$\|G : H\|$	the index of H in G
$\langle S \rangle$	the subgroup generated by S, where S is a subgroup of some group
$\widetilde{G}$	the left-regular representation of a group G
$\vec{\Gamma}_c(G, S)$	the edge-colored Cayley digraph of a group G with respect to generating set S
$\vec{\Gamma}(G, S)$	the Cayley digraph of a group G with respect to generating set S
$\Gamma(G, S)$	the Cayley graph of a group G with respect to generating set S
$A(X)$	the automorphism group of a (di)graph X
$Aut(G)$	the automorphism group of a group G
$\chi(X)$	the chromatic number of a graph X
$Orb_G(a)$	the orbit of a under a specified action by G
G_a	the stabilizer of a under a specified action by G
$X \Box Y$	the Cartesian product of (di)graphs X and Y
$H \times G$	the direct product of the groups H and G
$H \ltimes_\theta N$	the semi-direct product of N by H with respect to $\theta : H \to Aut(N)$
K_n	the complete graph on n vertices
1	the trivial group
$\mathbb{Z}_n$	the additive group of integers modulo n
D_n	the dihedral group with $2n$ elements
S_A	in Chapters 1, 3, and 4, the symmetric group of the set A; in Chapter 2, a generating set of a group A
S_n	the symmetric group on $\{1, \ldots, n\}$
A_n	the alternating group on $\{1, \ldots, n\}$
Q	the quaternion group with 8 elements

Abstract

This thesis is an exploration of the relationship between groups and their Cayley graphs. Roughly speaking, a group is a set of objects with a rule of combination. Given any two elements of the group, the rule yields another group element, which depends on the two elements chosen. A familiar example of a group is the set of integers with addition as the combination rule. Addition illustrates some of the properties which a group combination rule must have, including that it is associative and that there is an element that, like 0, doesn't change any element when combined with it. The information in a group can be represented by a graph, which is a collection of points, called "vertices," and lines between them, called "edges." In the case of the graph encoding a group, the vertices are elements of the group and the edges are determined by the combination rule. This graph is called a Cayley graph of the group.

We discuss the basic correspondences between the structure of a group and the structure of its Cayley graphs, focusing particularly on the group properties which can be "read" off of some of its Cayley graphs. We then examine Cayley graphs in the context of the graph regular representation question, providing a review of graph regular representation results and some of the more salient strategies used to obtain them. Finally, we present new results about compatible colorings of Cayley graphs, which are proper colorings such that a particular subgroup of the automorphism group permutes the color classes. We prove that a minimal Cayley graph does not have a compatible coloring if and only if it is the Cayley graph of a cyclic group whose order is a product of distinct primes with a specific generating set.

Acknowledgments

To my thesis advisor, Bill, for acting as a mathematically sophisticated ear throughout my exploration of a topic outside of his usual area of interest, and for his careful attention to the details of my exposition and proofs.

To my readers, Aaron and Pat, for their steady interest and encouragement during the thesis process.

To my professors, Péter Hermann and Gábor Simonyi, in Budapest, for introducing me to the beauties of groups and graphs and planting the seeds for the topic of this thesis.

To my research mentors, Edward Aboufadel, András Gyárfás, and Glencora Borradaile, for exposing me to mathematical research in a variety of guises and giving me the courage to embark on research independently.

To Budapest Semesters in Mathematics, and the REUs at Grand Valley State University, ICERM, and Oregon State University, for the deeply influential role they played in my development as a student of mathematics.

Finally, to the mathematics faculty and students of Simon's Rock, past and present, for the welcoming and rich community they form, both as scholars and as people; for their curiosity in the face of difficult ideas; and for their sincere willingness to teach and learn from each other, inside and outside of the classroom.

Contents

To the Reader

When one writes about mathematics, or any other field rife with specialized vocabulary, one faces something of a crisis of communication. The language used to discuss mathematical ideas is extraordinarily precise and unambiguous. It is also largely impossible to decode if you are unfamiliar with it, and sometimes masks, rather than elucidates, the intuition behind the concepts it describes. This thesis attempts to balance mathematical rigor with accessibility to a (fairly) general audience, presenting both an intuitive and a precise picture of the topics at hand. The focus is not on presenting new results (though these do form the fourth chapter), but rather on developing a multi-faceted understanding of the relationship between two kinds of mathematical objects, groups and graphs, and some of the interesting questions involving this relationship. The intended audience of this work is the student wishing to explore, not the expert seeking to prove. It should be comprehensible to anyone with some knowledge of group theory (about an undergraduate course worth), a bit of general experience with proofs, and a willingness to remember definitions. My assumption that my readers have this background in mathematics is my concession to the power of mathematical language; without words like *coset* and *normal subgroup*, it would be impossible to discuss the more complex, and more interesting, aspects of groups and their graphs.

That being said, this section is meant to give some idea of why I wrote the thesis I did— the origins of my interest in the topic and the way that interest developed into the work you are reading— in addition to what precisely that thesis consists of, and so should be accessible to everyone. We begin with "what," in the form of a description of the two mathematical structures at the center of the thesis. Roughly speaking, a group is a collection of objects with a rule of combination. If you take any two elements of a group, this rule will give you another element of the group. A familiar example is the integers $\mathbb{Z} := \{\ldots, -3, -2, -1, 0, 1, 2, 3, \ldots\}$, where the combination rule is normal addition. The combination rule must satisfy a number of axioms, which ensure it acts something like addition of integers— for example, there's always a group element that, like 0, leaves every element unchanged when combined with it. Groups are quite basic, in the sense that many more complicated mathematical objects have a group at their foundation (not in the sense that they are completely understood). They are also quite general. The group axioms are relatively unrestrictive, so all sorts of collections of objects with wildly different combination rules are groups.

The graphs we deal with here aren't graphs of a function in the plane. Instead, a graph is a collection of points (called "vertices") and some lines connecting them (called "edges"). Given a list of a graph's vertices and edges, it's quite easy to draw the graph on a piece of paper and, indeed, this is usually how we think of graphs. In this thesis, we are interested in graphs for their ability to encapsulate relationships between objects. For example, if you were throwing a party, you could make a graph where the vertices are the people you invited, and there's an edge between two vertices if those people know each other. This graph would be a visual representation of how well-acquainted the people invited to the party are as a group, and it helps you better understand your invite list. For instance, you could immediately see if there was one person who didn't know anyone else, as there would be a vertex with no edges in your graph.

I was introduced to group theory and graph theory during my junior year, which

I spent studying mathematics in Budapest. Towards the end of my second semester, my group theory professor mentioned a way to make a graph out of a group. The vertices of this graph are group elements and edges are based on the relationships the combination rule establishes. To be a bit more specific, to make the edges, we pick some elements from the group. Two group elements a and b are related (and so we put an edge between them) if a combined with one of the group elements we chose gives us b. Graphs created in this way from groups are called Cayley graphs; each group has a variety of Cayley graphs, whose structure depends on the elements we pick from the group in the beginning. But encoded in each Cayley graph are the distinguishing aspects of the group it came from, namely its elements and its combination rule. As a result, a Cayley graph is a visual representation of the group it came from, just as the graph we discussed earlier is a representation of the mutual well-aquaintedness of the people you invite to your party.

In class, Cayley graphs were discussed as part of a proof. They allow one to think of groups in a concrete, visual way and, as you might imagine, this is sometimes useful when you're trying to prove a statement about a group's properties. This is what initially drew me to Cayley graphs. I was fascinated by the translation from abstract to visual, and what was gained (and lost) in the process. Some group properties are visibly apparent in a group's Cayley graphs, but some are indiscernible. Different Cayley graphs of the same group can give you extremely different information, revealing distinct aspects of the group's structure. Chapter 2 explores the details of this translation, examining what the (algebraic) properties of a group imply about the (combinatorial) properties of its Cayley graphs. Chapter 3 continues along this vein, investigating a more subtle question about the relationship between groups and their Cayley graphs which was answered completely in the early 1980s.

True to the spirit of my liberal arts education, my interest in Cayley graphs was motivated as much by my experiences with analyzing texts as it was by my study of mathematics. There is a reason I think of creating a Cayley graph as an act of translation; group theory and graph theory are extremely disparate branches of mathematics, each with its own language developed to deal with very different structures, and each focusing on problems and properties that don't seem to relate to the other. Thinking of a group in terms of its Cayley graph places it in a completely different context, providing an alternate way to conceptualize ideas with which one may already be familiar. This exercise is fruitful, both for one's own understanding of group theory, and for the field of group theory in general; the nuances of another language frequently illuminate notions one may have missed in the original.

As I began to learn more about Cayley graphs, I noted that this translation occurs primarily in one direction, from group to graph. Mathematicians have thought deeply about what information the structure of a group gives us about its Cayley graphs, but there has been relatively little consideration of what the purely graph theoretic properties of a Cayley graph tell us about the group it represents. This struck me as unsatisfying on principle. We should, after all, be able to translate both ways. Cayley graphs are not just representations of groups; they are also graphs. Understanding how their properties in this context (i.e., that of graph theory) reflect on the groups they represent seems an essential part of understanding the connection between groups and their Cayley graphs. Chapter 4 pursues this direction of inquiry, presenting results that translate a specific Cayley graph property into a property of the groups they encode.

Chapter 1

The Basics

This chapter includes the preliminary definitions and notions necessary for the later discussion of groups and their Cayley graphs. Some familiarity with basic group theory (roughly the topics seen in an introductory undergraduate course) is assumed; see [14] or [11] if any of the concepts mentioned without definition are mysterious to you. The graph theory section should, however, be self-contained. For a more thorough exploration of the graph theory terms defined here, see [10]. For the sake of conciseness, the propositions and theorems of this chapter are frequently given without formal or full proofs unless the result is particularly exotic or the proof particularly illuminating.

1.1 Group Theory

First, some notation. We'll denote the trivial group by 1. $\mathbb{Z}_n := \{\overline{0}, \ldots, \overline{n-1}\}$ is the additive group of integers modulo n. D_n is the dihedral group on $2n$ vertices; we denote (clockwise) rotation by r and reflection (about a vertical axis) by s. S_A is the symmetric group of the set A, though when $A = \{1, \ldots, n\}$, we write S_n instead. The notation for subgroups and normal subgroups is standard; note that $\lneq$ denotes proper subgroup. The identity element of a group G is denoted e_G, or e if the group is obvious from context. The only exception to this is if G is an additive group, in which case its identity element is $\overline{0}$. Finally, the automorphism group of a group G is denoted $Aut(G)$.

Let G be a group and $S \subseteq G$ a subset of G. $\langle S \rangle$ is the smallest subgroup of G that contains S. Equivalently, $\langle S \rangle$ is the intersection of all subgroups H of G that contain S. $\langle S \rangle$ consists exactly of all products of integer powers of elements of S, which are also called *words* in S. If $\langle S \rangle = G$,then we say S *generates* G and call S a *generating set* of G. To put it differently, S generates G if every element of G can be written as a word in S; if G is finite, this is the same as saying every element of G can be written as a product of elements of S. The elements of a generating set are called *generators.* A generating set S is *minimal* if for all $s \in S$, $\langle S \setminus s \rangle \neq G$. Minimal generating sets don't contain any "redundant" generators, meaning that no generator can be written as a word in the other generators. Note that, unlike bases in a vector space, two minimal generating sets of the same group do not necessarily have the same cardinality. For example, the two generating sets $\{(12), (13), (14)\}$ and $\{(12), (1234)\}$ of S_4 are both minimal. Words in any pair of permutations in the first set fix at least one number and thus no two permutations in that set generate S_4, while the fact that S_4 is not cyclic means that if a permutation is removed from the second set, it no longer generates S_4.

We often are concerned with generating sets as they feature in group presentations. Informally, a presentation of a group G is a set S of generators and a description of how the generators relate to each other. This description is given by a set R of words in S that equal the identity, called *relators.* We write a group presentation $\langle S|R \rangle$, and call this a presentation of G *in terms of* S. If a group G is given in terms of a group presentation, we understand G to be the group generated by S whose elements are subject only to the relations in R. For example, a presentation of D_n is $\langle r, s|r^n, s^2, rsrs \rangle$. Minimal generating sets are particularly useful in the context of group presentations, as non-minimal generating sets involve "nonessential" relators which give us little or no new information about the structure of the group.

The intuitive notion of a group presentation given above will suffice for most of our discussion; those interested in a more formal definition should see [18]. For the sake of later clarity, we address one point in greater detail here. In a group presentation $\langle S|R\rangle$ of a group G, we think of every group element as a word in S. This is possible because S is a generating set. However, given $g \in G$, there is usually more than one word in S equal to g. For example, the identity element is equal to every relator. Similarly, if we have one word in S equal to g, we can obtain another by inserting a relator at the beginning or end of the word, or between two consecutive "letters." So to each element $g \in G$ there corresponds not just a single word in S, but an equivalence class of words. (The set of equivalence classes of words is in fact a group with respect to concatenation, which is isomorphic to G.) In some circumstances, particularly in Chapter 4, it is convenient to choose a unique representative from each equivalence class, which we can do by the process outlined in [18, Sec. 1.3].

This thesis focuses on the connections between a group and a graph which in some sense represents it. One way the group relates to the graph is through group action, so we recall some facts of group actions.

Definition 1.1. Let G be a group and A a set. A *left action of G on A* is an operation $* : G \times A \to A$ such that for all $a \in A$ and $g, h \in G$,

i) $e * a = a$;

ii) $g * (h * a) = (gh) * a$.

For example, we can define a very natural left action $*$ of D_5 on the vertices of the pentagon, labeled 1 through 5 clockwise with 1 at the top of the pentagon. Condition ii) of the definition of action guarantees that defining the action of r and s on the vertices completely determines how all other elements of D_5 act on the vertices, since

r and s generate the whole group. So we define $r * i := i + 1$ for $i \in \{1, \ldots, 4\}$, $r * 5 := 1$, $s * 1 := 1$, $s * 2 := 5$ and $s * 3 := 4$. In words, this action tells us how the vertices of a pentagon are shuffled when we apply different symmetries, represented by elements of D_5.

We can also think of a group action as a homomorphism from G to S_A, the symmetric group of A. Given a group action $*$ and a fixed $g \in G$, we associate to g the permutation γ_g that describes how g acts on A. $\gamma_g : A \to A$ is defined by $\gamma_g(a) := g * a$. It is indeed a permutation on A, since by the definition, its inverse is $\gamma_{g^{-1}}$. In the example given above of D_5 acting on a pentagon's vertices, $\gamma_r = (12345)$ and $\gamma_s = (25)(34)$. By condition ii), for $g, h \in G$, $\gamma_g \circ \gamma_h = \gamma_{gh}$, which means the map $g \mapsto \gamma_g$ is a homomorphism from G to S_A. Note that the kernel of this homomorphism is precisely those set of the elements of G that fix every element of A. We switch between speaking of a group action in terms of the permutations γ_g and in terms of the operation $*$ as is convenient.

Given a left action $*$ of G on A and an element $a \in A$, the *orbit* of a is $Orb_G(a) := \{g * a : g \in G\} = \{\gamma_g(a) : g \in G\}$. If the intersection of $Orb_G(a)$ and $Orb_G(b)$ is nonempty, then $Orb_G(a) = Orb_G(b)$; essentially, if γ_g takes a to b, then a can be mapped to the image of b under γ_h via composition. In the pentagon example, $Orb_{D_5}(1) = \{1, \ldots, 5\}$, as r^i takes 1 to $i+1$ for $i \in \{1, \ldots, 4\}$. This means that every vertex has the same orbit under the action of D_5 on the pentagon. The *stabilizer* of $a \in A$ is $G_a := \{g \in G : g * a = a\}$, the set of all group elements which fix a. It can easily be checked that G_a is a subgroup of G. The stabilizer of 1 under the action of D_5, for example, is $\{e, s\}$. The orbit and stabilizer of $a \in A$ are connected to each other in a number of ways. First, if two elements of A are in the same orbit, their stabilizers are conjugate; to be more precise, $G_{g*a} = gG_ag^{-1}$. This fact follows easily from the definitions. In the case of the action of D_5 on the pentagon, we mentioned that any vertex could be reached from 1 by powers of r. Thus, the stabilizer of any vertex is a conjugate of $\{e, s\}$ by some power of r. Second, there is a bijective correspondence between the left cosets of G_a and the elements of $Orb_G(a)$, given by $hG_a \mapsto h * a$. This gives rise to the following proposition.

Proposition 1.1 (Orbit-Stabilizer Lemma [9, p. 41])**.** *Let G be a group acting on a set A. Then for all $a \in A$, $|G : G_a| = |Orb_G(a)|$.*

If G is finite, the Orbit-Stabilizer Lemma implies $|Orb_G(a)|$ divides $|G|$; in particular, $|G| = |G_a||Orb_G(a)|$.

Group actions with certain properties give some insight to the structure of the group. An action is called *faithful* if the kernel of the associated homomorphism is trivial, which happens if and only if the intersection of all stabilizers is trivial. In this case, G is isomorphic to a subgroup of S_A. It is not difficult to verify that the action of D_5 on the pentagon is faithful, so $D_5 \cong \langle (12345), (25)(34) \rangle$. An action is *transitive* if $Orb_G(a) = A$ for some (and thus, all) $a \in A$. All stabilizers of a transitive action are conjugate subgroups of index $|A|$, as we discussed above in the case of the action of D_5 on the pentagon. If an action is transitive and has the property that all stabilizers are trivial, it is called *regular*. Regular actions are particularly easy to understand, as an action is regular if and only if for all $a, b \in A$, there is a unique $g \in G$ such that $g * a = b$. Note that regular actions are also faithful, so if G acts regularly on A, then G is a permutation group. Further, by the Orbit-Stabilizer Lemma, $|G| = |A|$.

One group action that will feature centrally in later discussion is the action of a group G on itself by left multiplication. To be more formal, this action is given by $* : G \times G \to G$, where $g * h := gh$, or, equivalently, by the $G \to S_G$ homomorphism $g \mapsto \lambda_g$, where $\lambda_g(h) := gh$. This action is regular, as given $f, h \in G$, $g = hf^{-1}$ is the unique element of G such that $g * f = h$. As mentioned above, this implies that the correspondence $g \mapsto \lambda_g$ is an isomorphism between G and a subgroup $\widetilde{G}$ of S_G. $\widetilde{G}$ is called the *left regular representation* of G; in a mild abuse of notation, we sometimes write $G \leq S_G$ or something to that effect, though of course it is actually the left regular representation of G that is a subgroup of S_G.

1.2 Graph Theory

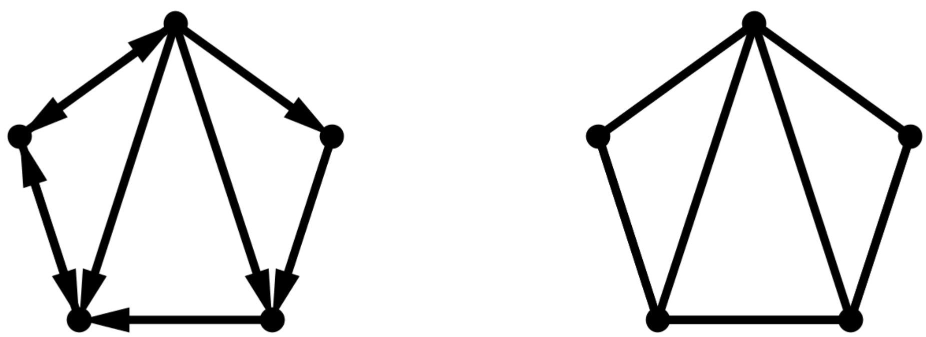

Figure 1.1: A directed graph and an undirected graph with the same underlying vertex set. If a directed graph has edges (v, w) and (w, v), we draw a single line from v to w with an arrow at both ends, as above.

We will consider two different kinds of graphs, directed and undirected. A *directed graph* $\vec{X} = (V(\vec{X}), E(\vec{X}))$ is a set $V(\vec{X})$ of *vertices* and a set $E(\vec{X})$ of ordered pairs of vertices, called *edges*. We denote an edge of a directed graph by (v, w), and say that there is an edge *from v to w*. Directed graphs can easily be represented visually, as in Figure 1.1. Vertices are drawn as points (sometimes labeled) and an edge (v, w) is an arrow pointing from vertex v to vertex w. An *undirected graph* $X = (V(X), E(X))$ is defined in much the same way: $V(X)$ is a set of vertices and $E(X)$ is a set of edges, which are here unordered pairs of vertices. An edge of an undirected graph is denoted $\{v, w\}$, and is drawn as a line segment between vertices. For both directed and undirected graphs, if a vertex v is in an edge with vertex w, we say v and w are *adjacent*, or w is a *neighbor* of v (and vice-versa). In order to compare the edge sets of directed and undirected graphs, we define the *underlying edge set* of a directed graph to be the edge set with directions removed; in other words, all ordered pairs are replaced by sets. Note that the underlying edge set of an undirected graph is just its edge set. A *subgraph* $X' = (V(X'), E(X'))$ of a graph $X = (V(X), E(X))$ is a graph where $V(X') \subseteq V(X)$ and $E(X') \subseteq E(X)$. If X' is a subgraph of X, we write $X' \subseteq X$.

In this thesis, all graphs are assumed to be *simple*, that is, without loops (edges consisting of a single vertex) or parallel edges (multiple occurrences of the same edge). Vertex sets are assumed to be finite. We may denote the vertex and edge sets of a graph simply as V and E, if the graph is clear from context. For convenience, we will from this point on refer to directed graphs as "digraphs" and undirected graphs as simply "graphs"; if we don't wish to specify directed or undirected, we will use "(di)graph."

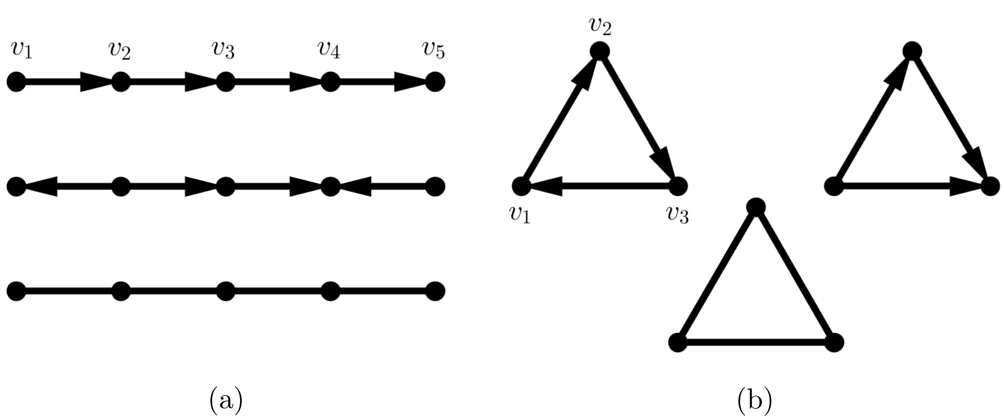

Figure 1.2: (a) Paths on 5 vertices, oriented, directed, and undirected. (b) Cycles on 3 vertices, oriented, directed but not oriented, and undirected.

Two families of (di)graphs will be of particular importance in later discussion: paths and cycles. A *path* on n vertices is a (di)graph with vertex set $\{v_1, \ldots, v_n\}$, where v_i and v_j are adjacent if $\{i, j\} = \{k, k+1\}$ for some $k \in \{1, \ldots, n-1\}$ (see Figure 1.2a). A cycle on n vertices is a path with one additional edge between v_1 and v_n (see Figure 1.2b). We say a path is *oriented* if all edges are directed from v_i to v_{i+1}; similarly, an *oriented cycle* is an oriented path with the additional edge (v_n, v_1). We'll use path and cycle to speak of both directed and undirected graphs; it should be clear from context which we're referring to.

At the most basic level, a (di)graph is a set with some special structure. There are many mathematical objects of this sort, among them groups, a set with a particular sort of operation, and topological spaces, a set with distinguished subsets. For (di)graphs, the structure on the underlying set of vertices is a binary relation "is adjacent to," or, in other words, the edges. As with groups and topological spaces, one way to understand (di)graphs is to study functions on the vertices that preserve this additional structure.

Definition 1.2. Let $\vec{X} = (V(\vec{X}), E(\vec{X}))$ and $\vec{Y} = (V(\vec{Y}), E(\vec{Y}))$ be digraphs. A function $\alpha : V(\vec{X}) \to V(\vec{Y})$ is a *digraph homomorphism* from $\vec{X}$ to $\vec{Y}$ if for all $v, w \in V(\vec{X})$, $(v, w) \in E(\vec{X})$ implies $(\alpha(v), \alpha(w)) \in E(\vec{Y})$.

It's clear from the definition that digraph homomorphisms also preserve the direction of edges. Graph homomorphisms are defined similarly as functions from the vertices of one graph to those of another that preserve edges, though of course with graphs this no longer involves preserving edge direction. From the perspective of homomorphisms, there is a natural correspondence between graphs and digraphs with a symmetric edge set (i.e. if $(v, w) \in E$, then $(w, v) \in E$). We associate to each graph X a digraph $\widehat{X}$ on the same vertex set, where $(v, w) \in E(\widehat{X})$ if and only if $\{v, w\} \in E(X)$. A graph homomorphism α from X to Y is also a digraph homomorphism from $\widehat{X}$ to $\widehat{Y}$. Similarly, a digraph homomorphism between two digraphs with symmetric edge sets is also a graph homomorphism between the corresponding graphs. For this reason, we limit the following discussion to digraphs, with the understanding that the same statements also hold for graphs because of this correspondence.

Roughly speaking, a digraph homomorphism from $\vec{X}$ to $\vec{Y}$ takes $\vec{X}$ and "fits" it into $\vec{Y}$ so that its edges match up with (some of) $\vec{Y}$'s edges (see Figure 1.3). Some amount of deformation is allowed in this process: non-adjacent vertices of $\vec{X}$ could be placed on the same vertex of $\vec{Y}$, for example. It's also important to note that homomorphisms do not preserve non-adjacency, so even if two vertices of $\vec{Y}$ aren't

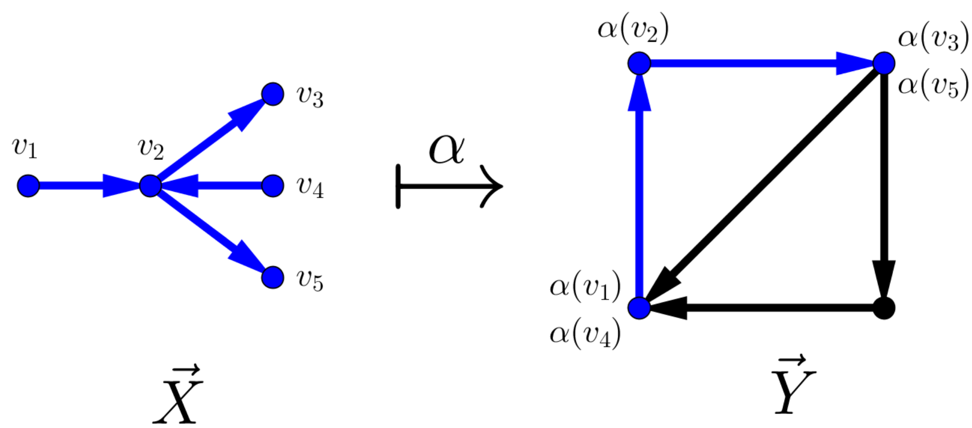

Figure 1.3: A digraph $\vec{X}$ and its image under a homomorphism to $\vec{Y}$, shown in blue.

adjacent, their images under a homomorphism might be.

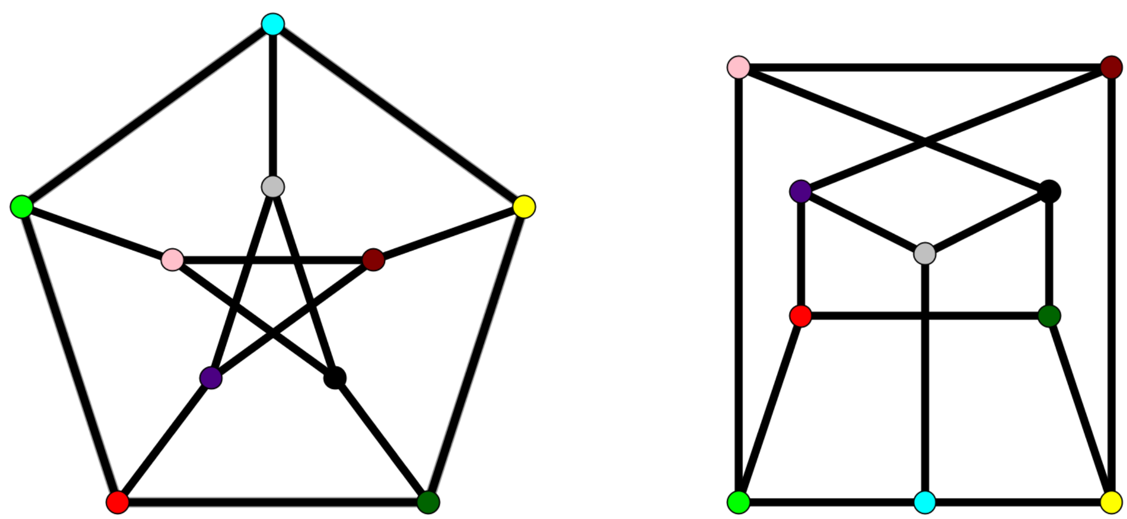

Figure 1.4: Two isomorphic graphs. An isomorphism between the two is given by mapping vertices in the first to vertices of the same color in the second.

As one might anticipate, a bijective digraph homomorphism whose inverse is also a homomorphism is a *digraph isomorphism*. Digraph isomorphisms do preserve non-adjacency, a property which follows fairly immediately from their definition. In fact, we could define a digraph isomorphism as a bijective function $\alpha : V(\vec{X}) \to V(\vec{Y})$ such that $(v, w) \in E(\vec{X})$ if and only if $(\alpha(v), \alpha(w)) \in E(\vec{Y})$; the two definitions are equivalent. If two digraphs are isomorphic, they have the same underlying structure even if they appear to be different. This statement has a very nice visual interpretation. Say you have two digraphs drawn on a piece of paper. If they're isomorphic, you can take one of the digraphs and place its vertices on top of the vertices of the other in such a way that the edges of the two digraphs exactly match up. The isomorphism tells you how to place the vertices of the first digraph so that this happens. Figure 1.4 gives an illustration of this process.

A digraph isomorphism from $\vec{X}$ to itself is called a *digraph automorphism*. Automorphisms are symmetries of the digraph, or transformations of the digraph that result in a digraph "identical" to the first, just as reflecting a square across one of its diagonals gives a square that looks exactly the same (though the locations of some vertices have changed). Our previous intuition about isomorphisms applies to automorphisms as well: if we think of a digraph drawn twice on a piece of paper, an automorphism will tell us how to take the vertices from the first drawing and place them on the vertices of the second drawing so the edges match up. There is one very easy way to do this for any digraph, which you have perhaps already thought of. If you don't rearrange the vertices at all, and simply place each vertex in the first drawing on its counterpart in the second, the edges of the two will certainly match. This automorphism, described by the mapping $v \mapsto v$ for all $v \in V$, is called the trivial automorphism. For some (in fact, most) digraphs, the trivial automorphism is the only automorphism. We call such digraphs *asymmetric*. The oriented path on n vertices is an example. Roughly speaking, asymmetric graphs have vertices that can be structurally distinguished from every other vertex; perhaps they have a different number of neighbors, or a different number of edges going out of them or coming in to them. These distinguishable vertices must be fixed by any automorphism, which, because of their distribution in the digraph, forces all of the other vertices to be fixed as well. The oriented path on n vertices provides an example of this. The vertices at the beginning and end of the path are distinguishable from every other vertex, as the former is the only vertex with no edge going in and the latter is the only vertex with no edge going out. So every automorphism of the path fixes the beginning and end vertices, and are forced by adjacency preservation to fix all other vertices as well.

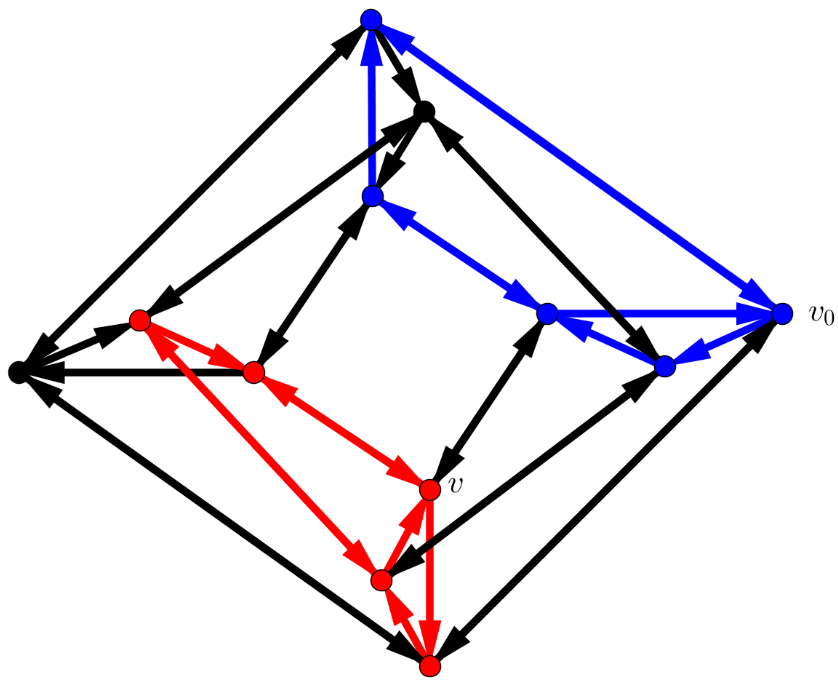

Figure 1.5: For any vertex v of this digraph, you can find a subgraph $\vec{X}$ isomorphic to the blue subgraph such that v has the same position in $\vec{X}$ as v_0 has in the blue subgraph. An example is shown in red.

In this thesis, we're more interested in (di)graphs with many automorphisms than in those with none. Intuitively, these (di)graphs are fairly uniform, in the sense that they have many vertices that "look" the same. These vertices have the same number of neighbors and have the same number of edges coming in to and going out of them; more significantly, if one is in some subgraph $\vec{Y}$, every other is in a subgraph isomorphic to $\vec{Y}$ in the same "place" (see Figure 1.5). These vertices are structurally indistinguishable, and so one could be the image of another under an automorphism. It's important to note that this is a necessary, not sufficient, condition for the existence of a nontrivial automorphism, as symmetry is a global property, not a

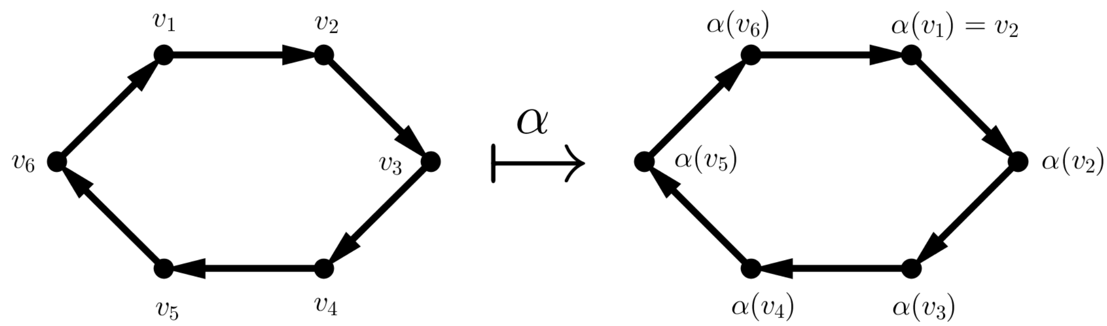

(a) α, an automorphism of the oriented cycle on 6 vertices, is a "rotation" clockwise by one vertex. All automorphisms of oriented cycles are in fact rotations clockwise by some number of vertices: if v_1 maps to v_k, then v_2 must map to v_{k+1} to preserve the (v_1, v_2) edge. Proceeding by this line of reasoning yields that the automorphism is a clockwise rotation of each vertex by k vertices.

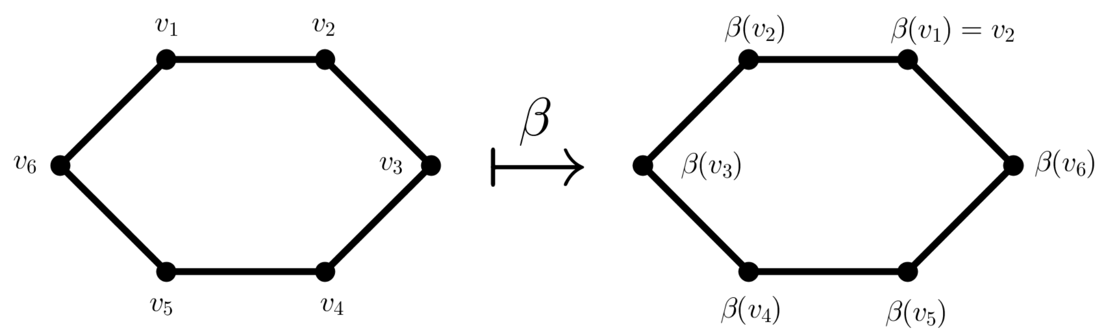

(b) All automorphisms of directed cycles are also automorphisms of undirected cycles. Undirected cycles also have an additional kind of automorphism, shown here. β is a reflection of the vertices across the vertical axis. All automorphisms of undirected cycles correspond to a combination of rotations and reflections.

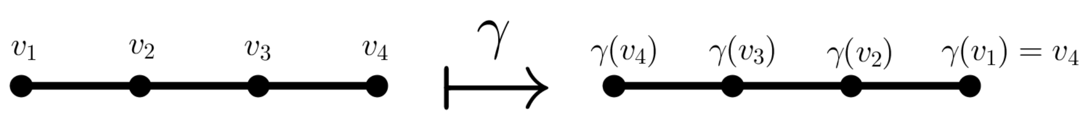

(c) Oriented (and most directed) paths have no nontrivial automorphisms, but undirected paths do have one, here called γ. The image of an end vertex under an automorphism must be an end vertex, because all other vertices have 2 neighbors. If you send v_1 to v_1, the rest of the vertices must be fixed and you get the trivial automorphism. If you send v_1 to v_n, then the image of the rest of the vertices is completely determined and the automorphism reverses the order of the vertices.

Figure 1.6

local one. Perhaps the least interesting example of a graph with many automorphisms is the complete graph on n vertices, denoted K_n, which has every possible edge (i.e. for all $v, w \in V(K_n)$, $\{v, w\} \in E(K_n)$). Any two vertices of K_n play exactly the same role in the graph, and, indeed, any bijective function $\alpha : V(K_n) \to V(K_n)$ is an automorphism. Oriented cycles and undirected cycles also have a number of automorphisms, which correspond to rotations of the vertices, while paths have only two, the trivial automorphism and the automorphism switching the vertices at the ends of the path (see Figure 1.6). Arbitrary directed paths and cycles rarely have nontrivial automorphisms, illustrating the fact that a digraph has at most as many automorphisms as a graph with the same underlying edge set, as edge direction adds an additional condition to preserving edges and thus places further constraints on automorphisms.

Given two automorphisms, it is intuitively plausible that their composition is also an automorphism; rearranging the vertices of a digraph $\vec{X}$ as directed by an automorphism, and then doing so again to the resulting digraph, should give a digraph isomorphic to $\vec{X}$, as adjacency and non-adjacency were preserved in each rearrangement. The fact that the automorphisms are from $\vec{X}$ to $\vec{X}$ would make their composition an automorphism rather than just an isomorphism. It's easy to check from the definition of automorphism that this intuition is correct, and the set of automorphisms of a given digraph is closed under composition. In fact, the automorphisms of a digraph $\vec{X}$ form a group with respect to composition, with the trivial automorphism as the identity. We denote this group $A(\vec{X})$.

There are a few important, if basic, things to note about the automorphism group of a digraph $\vec{X}$. First, since automorphisms are bijections with the same domain and codomain, $A(\vec{X})$ is a subgroup of S_V, the group of permutations of V. This means $A(\vec{X})$ inherits the natural action of S_V on the vertices of $\vec{X}$. This action is precisely what you would expect— a permutation α acts on a vertex by sending it to its image under α. More precisely, the action of $A(\vec{X})$ on V is given by $* : A(\vec{X})\times V \to V$, where for $v \in V, \alpha \in A(\vec{X})$, $\alpha * v = \alpha(v)$. This operation clearly satisfies the conditions for group action, as for $v \in V, \alpha, \beta \in A(\vec{X})$, $\beta * (\alpha * v) = \beta(\alpha(v)) = (\beta \circ \alpha) * v$. The identity condition is clear. This action is always faithful. It also allows us to use the Orbit-Stabilizer Lemma (Proposition 1.1) to find the size of $A(\vec{X})$, which is occasionally useful. The orbit of a vertex v consists of all images of v under an automorphism, while the stabilizer of v is all automorphisms which fix v.

A digraph whose automorphism group acts transitively on its vertices is called *vertex-transitive*. Vertex-transitive digraphs are very symmetric; in the sense of the earlier discussion about symmetric digraphs, all of their vertices "look" identical and play the same structural role in the digraph. By the Orbit-Stabilizer Lemma, $|V(\vec{X})|$ divides $|A(\vec{X})|$ for vertex-transitive $\vec{X}$, a fact we will use later.

The final graph-theoretic concept we need is that of colorings. A *coloring* of the vertices of a (di)graph X is a function $c : V \to \{c_1, \dots, c_n\}$, where $\{c_1, \dots, c_n\}$ is a set of colors (usually integers). It is sometimes more convenient to think of a coloring as a partition of the vertices rather than as a function. In this case, we define color classes $C_{c_i} := \{v \in V : c(v) = c_i\}$ and describe the coloring by giving the set of color classes, $C := \{C_{c_1}, \dots, C_{c_n}\}$. A vertex coloring c is called *proper* if for any adjacent $v, w \in V$, $c(v) \neq c(w)$. The color classes of a proper vertex coloring are an example of independent sets, which are subsets of the vertex set containing no adjacent vertices. The *chromatic number* of a (di)graph X, denoted $\chi(X)$, is the smallest integer m such that X has a proper vertex coloring using m colors.

Edge colorings are defined similarly as a function from E to a set of colors. Though many vertex-coloring concepts have edge-coloring analogues, we have need only of the vertex-coloring versions, so we refer interested readers to [10, Sec. 5.3].

1.3 Cayley Graphs

We now define the objects on which the remainder of the thesis focuses. From this point forward, we restrict our attention to finite groups, though many of the notions discussed below also hold for infinite groups.

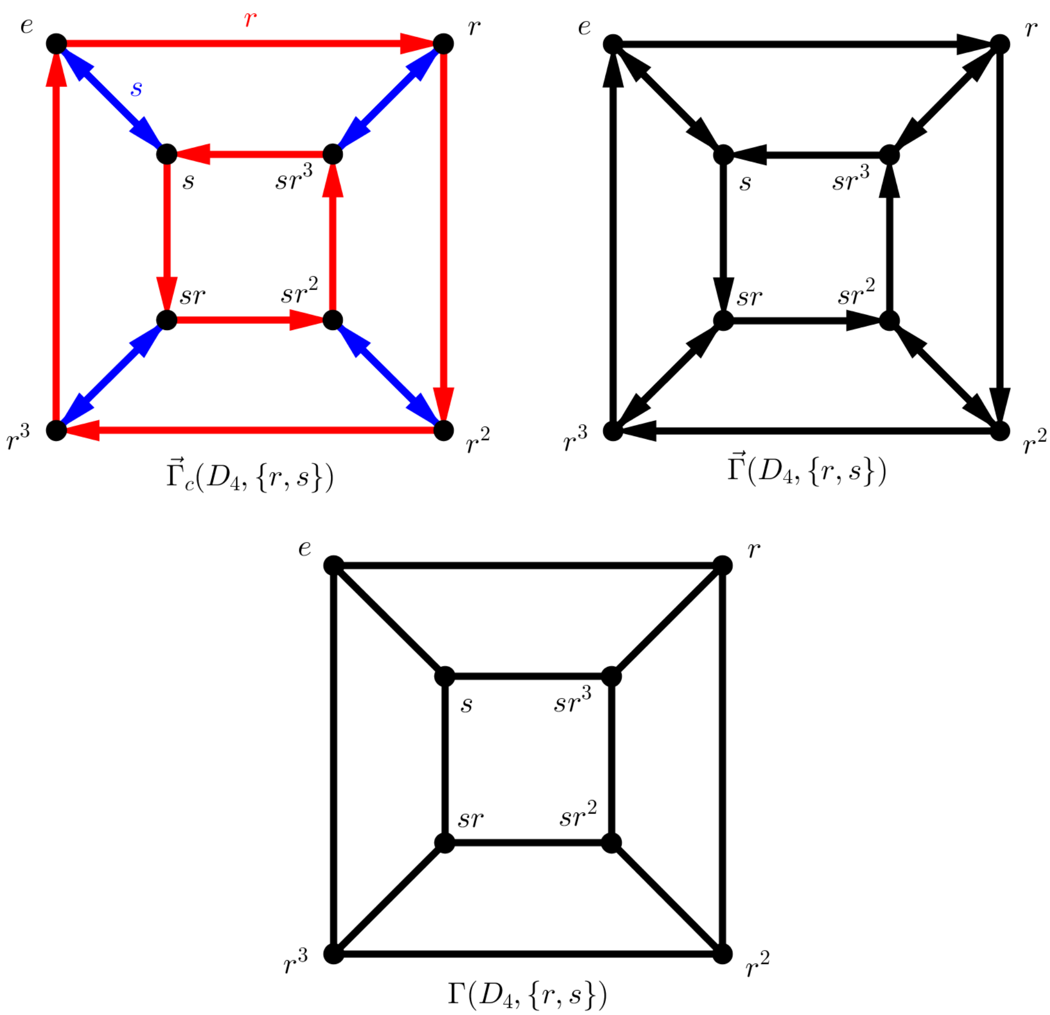

Figure 1.7: The Cayley (di)graphs of D_4, with $S = \{r, s\}$.

Definition 1.3. Let G be a group and $S \subseteq G$ a generating set of G. The *edge-colored Cayley digraph* of G with respect to S, denoted $\vec{\Gamma}_c(G, S)$, has vertex set $V = G$ and edge set $E = \{(g, gs) : g \in G \text{ and } s \in S\}$. The edges are colored by $j : E \to S$, where $j(g, gs) = s$. G is called the *generating group* of $\vec{\Gamma}_c(G, S)$.

As we are limiting our attention to graphs without loops, we assume in this definition that S does not contain e_G. An equivalent, if slightly less explicit, way to define the edges of $\vec{\Gamma}_c(G, S)$ are those ordered pairs of vertices (g, h) such that $g^{-1}h \in S$; the edge (g, h) is then colored $g^{-1}h$. From this alternate definition, it is easy to see that if S contains s and its inverse s^{-1}, then the edges colored s and the edges colored s^{-1} are between the same vertices, but in the opposite direction. More specifically, s-colored edges point from g to gs, while s^{-1}-colored edges point from gs to g.

If the generating set S is minimal, we call $\vec{\Gamma}_c(G, S)$ *minimal*. From the definition, it's easy to see that if S and S' are two generating sets of a group G such that $S' \subseteq S$, $\vec{\Gamma}_c(G, S') \subseteq \vec{\Gamma}_c(G, S)$. This implies that any edge-colored Cayley digraph contains a minimal edge-colored Cayley digraph.

Two related (di)graphs are also of interest to us. The *Cayley digraph* of G with respect to S, denoted $\vec{\Gamma}(G, S)$, is defined in the same way as the edge-colored Cayley digraph, but without the edge-coloring j. The *Cayley graph* of G with respect to S, denoted $\Gamma(G, S)$, is $\vec{\Gamma}(G, S)$ with the edge directions removed; in other words, $E(\Gamma(G, S)) = \{\{g, gs\} : g \in G \text{ and } s \in S\}$. See Figure 1.7 for examples of all 3 kinds of Cayley (di)graphs. Since edge-colored Cayley digraphs are in some sense the closest to their generating group, and many of their properties are shared by the non-colored Cayley (di)graphs, we discuss their properties first.

Edge-colored Cayley digraphs give a visual map of how the generators of a group interact to give the full group. Their structure is very much dependent on the generating set chosen, as shown in Figure 1.8. It is also connected to the structure of their generating groups, a relationship we'll explore more fully in Chapter 2. Here, we mention some properties of edge-colored Cayley digraphs that are fairly immediate from their definition. They are *connected*, meaning that there is a (oriented) path between any two vertices. This follows immediately from S being a generating set of G and is in fact the reason we assume S is a generating set in the definition of $\vec{\Gamma}_c(G, S)$ (if S is

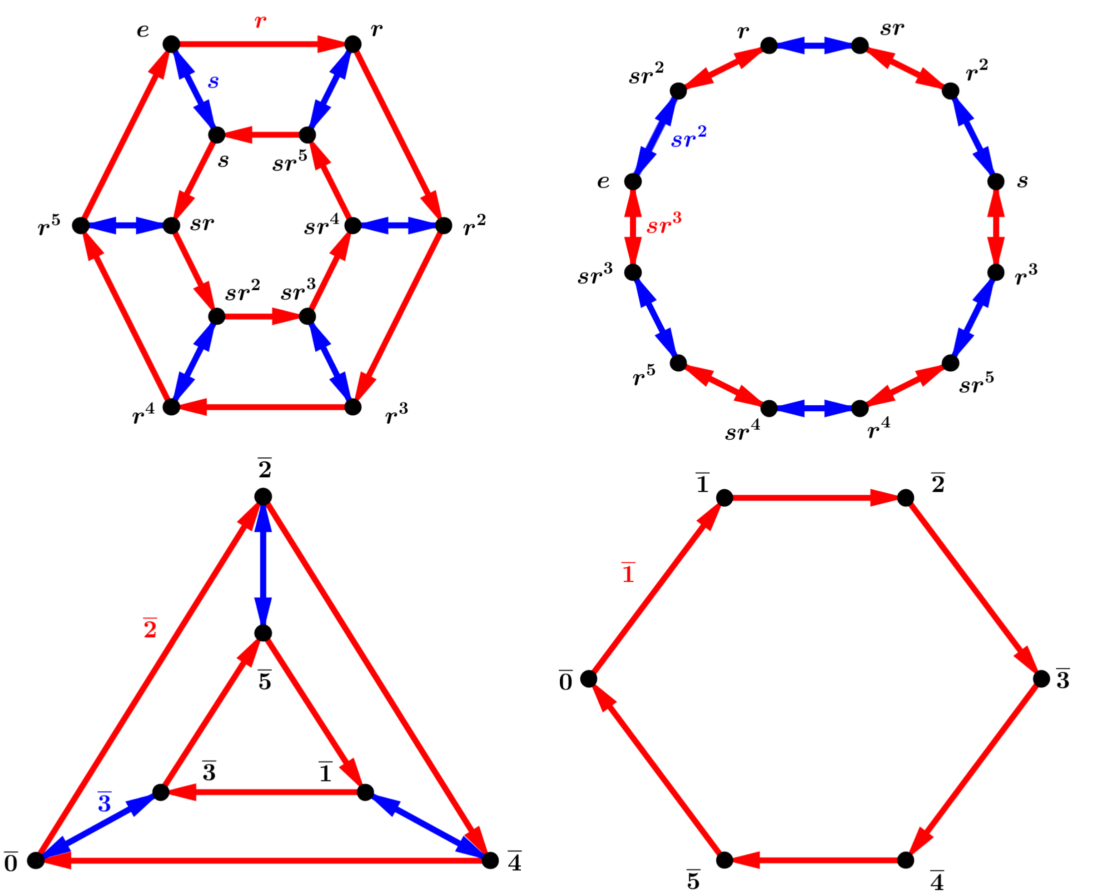

Figure 1.8: From the upper left corner, clockwise: $\vec{\Gamma}_c(D_6, \{r, s\})$, $\vec{\Gamma}_c(D_6, \{sr^2, sr^3\})$, $\vec{\Gamma}_c(\mathbb{Z}_6, \{\overline{1}\})$, $\vec{\Gamma}_c(\mathbb{Z}_6, \{\overline{2}, \overline{3}\})$. As you can see, the Cayley (di)graphs of a given group can be quite different, even if the generating sets used are minimal or of the same cardinality.

just an arbitrary subset of G, then the connected components of $\vec{\Gamma}_c(G, S)$ are cosets of $\langle S \rangle$). Additionally, every vertex g has exactly $|S|$ edges coming in to it, namely the edges $\{(gs^{-1}, g) : s \in S\}$. Clearly, there is exactly one edge of each color in this set. Similarly, each vertex has exactly $|S|$ edges going out of it, again one of each color. This sort of uniformity hints that edge-colored Cayley digraphs have nontrivial automorphisms, which is the case. Since edge color is an important feature of this digraph, we limit our attention to automorphisms which preserve the color of edges, which we denote in the usual manner by $A(\vec{\Gamma}_c(G, S))$.

The color-preserving automorphisms of $\vec{\Gamma}_c(G, S)$ are functions we have encountered before. For any $g \in G$, the left regular representation of g, $\lambda_g \in \widetilde{G}$, is such an automorphism. Indeed, for $h, k \in G$, there is an edge from $\lambda_g(h) = gh$ to $\lambda_g(k) = gk$ if and only if $(gh)^{-1}(gk) = h^{-1}k$ is in S, which in turn is true if and only if (h, k) is an edge. Edge color is also preserved, since $(gh, gk) \in E$ is colored $(gh)^{-1}(gk) = h^{-1}k$. So we have that $\widetilde{G} \subseteq A(\vec{\Gamma}_c(G, S))$. In fact, the following statement is true.

Proposition 1.2. *Let G be a group and $S \subseteq G$ a generating set. Then $\widetilde{G} = A(\vec{\Gamma}_c(G, S))$.*

Proof. ([7]) We showed above that $\widetilde{G} \leq A(\vec{\Gamma}_c(G, S))$.

To see that the opposite inclusion holds, let $\alpha \in A(\vec{\Gamma}_c(G, S))$. For $(h, hs) \in E(\vec{\Gamma}_c(G, S))$, $(\alpha(h), \alpha(hs))$ must also be an edge of $\vec{\Gamma}_c(G, S)$ colored s. This implies that $\alpha(hs) = \alpha(h)s$ for all $h \in G, s \in S$.

Given $g \in G$, we can write g as a product of elements of S, since G is finite. Say $g = s_{\nu_1} \cdots s_{\nu_n}$. Then

$$\begin{aligned}
\alpha(g) &= \alpha(eg) \\
&= \alpha(es_{\nu_1} \cdots s_{\nu_n}) \\
&= \alpha(es_{\nu_1} \cdots s_{\nu_{n-1}})s_{\nu_n} \\
&\vdots \\
&= \alpha(e)s_{\nu_1} \cdots s_{\nu_n} \\
&= \alpha(e)g.
\end{aligned}$$

Thus, $\alpha = \lambda_{\alpha(e)}$.

□

This proposition gives us some information about the Cayley (di)graphs of a group G. It implies that the action of color-preserving automorphisms on the vertices of $\vec{\Gamma}_c(G, S)$ is exactly the action of G on itself by left multiplication. As discussed in

Section 1, the latter action is regular, so the action of $A(\vec{\Gamma}_c(G,S))$ on $V(\vec{\Gamma}_c(G,S))$ is regular as well. In particular, $\vec{\Gamma}_c(G,S)$ is vertex-transitive. Since $A(\vec{\Gamma}_c(G,S)) \leq A(\vec{\Gamma}(G,S) \leq A(\Gamma(G,S)))$, the uncolored Cayley (di)graphs of a group are also vertex-transitive and their automorphism groups have a subgroup that acts regularly on their vertices. We will call such a subgroup "regular."

A quick note here about an implication of vertex-transitivity. We frequently omit vertex labels in Cayley (di)graphs, since by vertex-transitivity, we can label any vertex with the identity element. In an edge-colored Cayley digraph $\vec{\Gamma}_c(G,S)$, the vertex we label with the identity element completely determines the labels of the remaining vertices: say $v \in V(\vec{\Gamma}_c(G,S))$ has been labeled e. For $w \in V(\vec{\Gamma}_c(G,S))$, there is exactly one automorphism α taking v to w. $\alpha = \lambda_g$ for some $g \in G$, by what we've already discovered of the automorphisms of $\vec{\Gamma}_c(G,S)$, so w must be labeled g. In uncolored Cayley (di)graphs, there is more flexibility, since there are in general more automorphisms, but we can still give a valid labeling of a Cayley (di)graph based solely on the placement of the identity label.

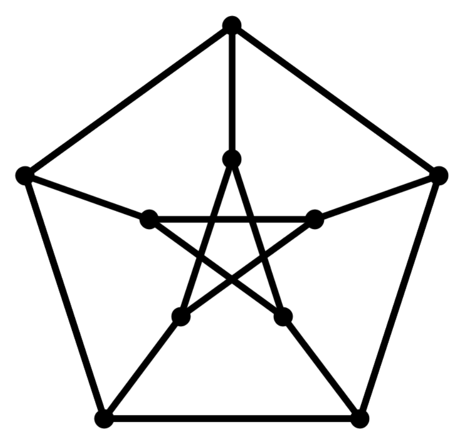

Figure 1.9: The Petersen graph is not a Cayley graph of any group, though it is vertex transitive.

It seems reasonable to ask if the properties of Cayley (di)graphs' automorphism groups mentioned above characterize Cayley (di)graphs. That is, is every vertex-transitive (di)graph a Cayley (di)graph for some group? Or is every (di)graph whose automorphism group has a regular subgroup a Cayley (di)graph? The answer to the first question is no; the graph in Figure 1.9 is a counterexample. Sabidussi answered the second question positively in [24] for graphs. We provide a slightly modified proof that gives the statement for both graphs and digraphs.

Proposition 1.3 ([24]). *Let X be a (di)graph. Then there exists a group G and $S \subseteq G$ such that $X = \Gamma(G, S)$ if and only if $A(X)$ contains a subgroup $\widetilde{G}$ acting regularly on $V(X)$.*

Proof. (Sketch) Suppose $\widetilde{G} \leq A(X)$ acts regularly on $V(X)$. Fix $v_0 \in V(X)$. For $u \in V$, there exists a unique $\alpha_u \in \widetilde{G}$ such that $\alpha_u(v_0) = u$, by the definition of a regular action. $\widetilde{G} = \{\alpha_u : u \in V(X)\}$, since the α_u's are distinct and $|\widetilde{G}| = |V(X)|$. Let $\{v_1, \ldots, v_n\}$ be the neighbors of v_0 if X is a graph; if X is a digraph, let $\{v_1, \ldots, v_n\}$ be the neighbors of v_0 such that (v_0, v_i) is an edge. Let $S := \{\alpha_{v_1}, \ldots, \alpha_{v_n}\}$. Then $\Phi : V(\Gamma(\widetilde{G}, S)) \to V(X)$ defined by $\Phi(\alpha_u) = \alpha_u(v_0)$ is an isomorphism. □

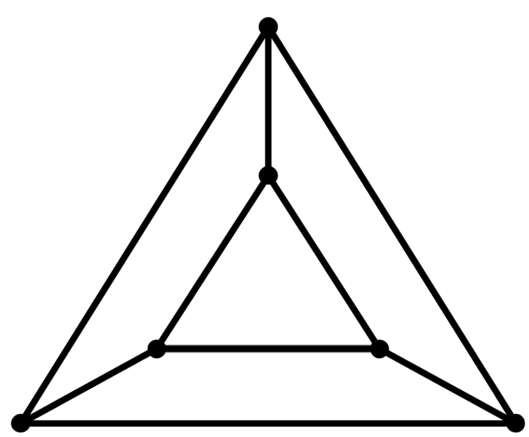

Figure 1.10: Nonisomorphic groups may have isomorphic Cayley graphs. The above graph, for example, is $\Gamma(\mathbb{Z}_3 \times \mathbb{Z}_3, \{(\overline{1}, \overline{0}), (\overline{0}, \overline{1})\})$, $\Gamma(D_3, \{r, s\})$, and $\Gamma(S_3, \{(123), (12)\})$.

So, a graph X is a Cayley graph of exactly the groups isomorphic to regular subgroups of its automorphism group. As an automorphism group may have multiple non-isomorphic regular subgroups, one graph may be the Cayley graph of a number of different groups. In other words, non-isomorphic groups can have isomorphic Cayley graphs. For example, the graph in Figure 1.10 is a Cayley graph of $\mathbb{Z}_3 \times \mathbb{Z}_3$, D_3 and S_3. Non-isomorphic groups cannot, of course, have isomorphic edge-colored Cayley digraphs, since isomorphic graphs have isomorphic automorphism groups; this provides some insight on how much information is lost when one passes from an edge-colored Cayley digraph to a Cayley graph with respect to the same generating set.

Chapter 2

Visualizing Groups Using Cayley Graphs

We begin our examination of Cayley (di)graphs by focusing on how these (di)graphs visually encapsulate their generating group, following Nathan Carter [5]. In contrast to the majority of this thesis, we consider only edge-colored Cayley digraphs in this chapter. These Cayley digraphs are, in some sense, the "closest" to their generating groups, as their automorphism groups are isomorphic to their generating groups (Proposition 1.2). It is also easier to (literally) see the algebraic information contained in edge-colored Cayley digraphs, so it seems reasonable to focus on them in particular when trying to understand the basic correspondences between a group and its Cayley (di)graphs. Minimal Cayley digraphs also have a privileged role in this chapter, as they are the easiest to understand from a visual, intuitive perspective. The additional information contained in a non-minimal Cayley (di)graph complicates the picture, and is generally redundant in this context.

As we'll see below, it's perhaps more accurate to say that a Cayley (di)graph reflects a certain presentation of its generating group (namely, the group's presentation in terms of the chosen generating set), rather than its generating group in general.

This is intuitively reasonable, as in a Cayley (di)graph, we understand the interactions of elements of the generating group through the interactions of the generators, which is exactly what a presentation describes. Of course, a group has many different presentations, and different presentations make different aspects of a group's structure more apparent. For example, if $G \cong \mathbb{Z}_6$, the presentation $G = \langle a | a^6 \rangle$ emphasizes that G is cyclic, while the presentation $G = \langle b, c | b^3, c^2, c^{-1}b^{-1}cb \rangle$ displays that G is a direct product of two cyclic groups. For this reason, different Cayley (di)graphs of the same group generally display different aspects of the group's structure. (Incidentally, this is also precisely why the Cayley (di)graphs of a given group can be extremely different, as shown in Figure 1.8.) A group property appears in every Cayley (di)graph of a group only if it is "expressed" in every group presentation, which is generally not the case.

2.1 Significant Subgraphs

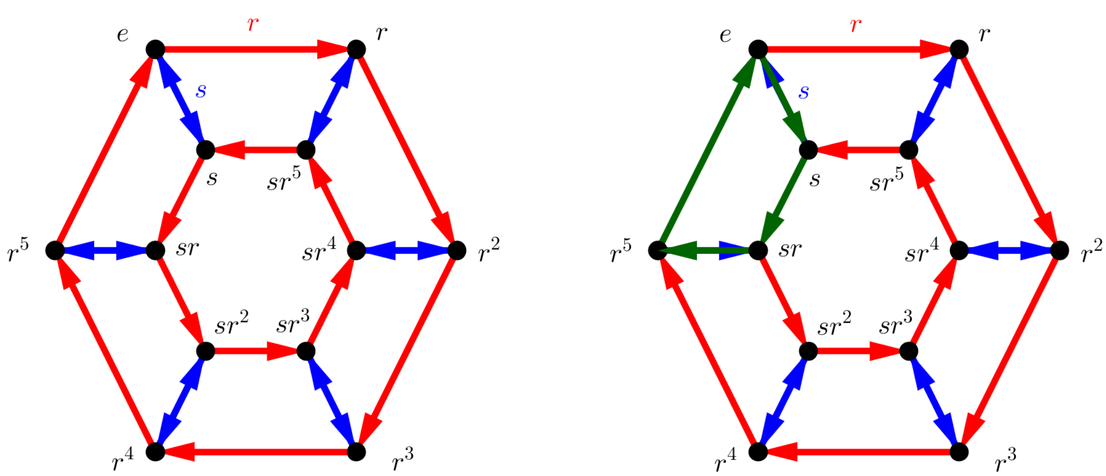

Figure 2.1: Two copies of $\vec{\Gamma}_c(D_6, \{r, s\})$. On the left, the red cycles reveal the order of the generator r. On the right, the green cycle (with 4 edges) demonstrates the order of sr.

We first consider certain subgraphs of $\vec{\Gamma}_c(G, S)$ and the information they yield about G and S. Oriented cycles in $\vec{\Gamma}_c(G, S)$ correspond to relations among generators:

since edges correspond exactly to right multiplication by a generator, if you can "walk" along some sequence of edges colored $s_{\nu_1}, \ldots, s_{\nu_n}$ and return to your starting place, the corresponding word $s_{\nu_1} \cdots s_{\nu_n}$ in S equals the identity. For example, in Figure 2.2, if you begin at any vertex and walk along two i-colored edges and then two j-colored edges, you end up back at the vertex you started, which means $i^2 j^2 = e$. In this way, you can read the presentation of G in terms of S off of its Cayley digraph $\vec{\Gamma}_c(G, S)$. If you choose a vertex v, the colors of edges going out from v will give you S and the oriented cycles in which v is a vertex will give you all essential relations between the generators. We note that in this process, you need not walk along an edge in the direction it points; walking along an s-colored edge (g, gs) backwards corresponds to right-multiplication by s^{-1}, since you're walking from gs to g. So one could walk along two i-colored edges in Figure 2.2 and then backwards along two j-colored edges to return to where you started, yielding the relation $i^2 j^{-2} = e$.

Monochromatic oriented cycles correspond to relations involving only one generator, i.e. a generator's order. The number of edges in a monochromatic cycle colored s is the order of s. The order of an arbitrary element can be seen in a Cayley graph in a similar, though less immediately apparent, fashion. Say g is a group element that can be expressed as a word of length n in the generators. This word corresponds to a sequence of edge colors. If you begin at a vertex and repeatedly walk along that particular sequence of colored edges until you return to where you started (see Figure 2.1), the number of edges in the cycle you walked divided by n is the order of g. Note that you can do this beginning at any vertex, since each vertex has edges of each color going in and out.

The connection between generator order and monochromatic cycles provides a nice visual confirmation of the fact that the order of group elements divides the order of the group. In the Cayley graphs displayed so far, you may have noticed that monochromatic cycles of a given color cover the vertices of the graph (see Figure 2.2).

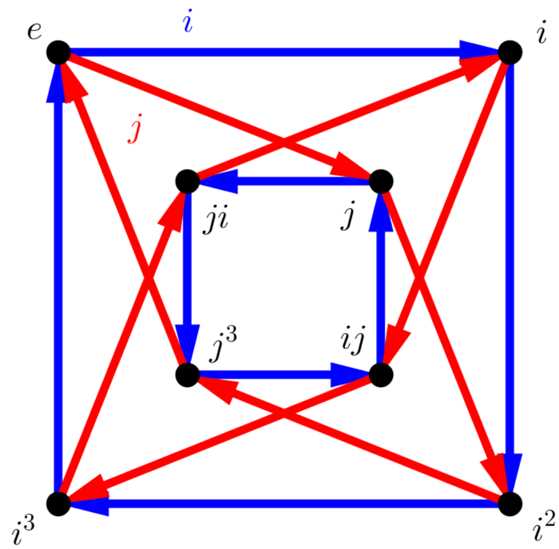

Figure 2.2: Monochromatic cycles partition the vertices of $\vec{\Gamma}_c(Q, \{i, j\})$, where $Q = \langle i, j | i^4, i^2 j^{-2}, j^{-1}ij = i^{-1}\rangle$ is the quaternion group.

Given a generator, every vertex is in a monochromatic cycle of that color; even better, the monochromatic cycles are disjoint, so every vertex is in only one. Since the length of these monochromatic cycles is precisely the order of the corresponding generator, if this partition of vertices by monochromatic cycles happens in general, it implies that the order of the generator divides the group order. This partition does indeed occur in every Cayley graph: if s is a generator and g an arbitrary group element, g is in precisely 2 s-colored edges, (gs^{-1}, g) and (g, gs). This means that in the subgraph of the Cayley graph containing all vertices and only s-colored edges, every vertex incident to precisely two edges. It's a basic fact from graph theory that such a graph is a disjoint union of cycles, so we conclude that every vertex of the Cayley graph is in precisely one s-colored cycle and thus the order of a generator divides the group order. Since every element is a generator in $\vec{\Gamma}_c(G, G)$, this implies that the order of any group element must divide the group order.

A number of other relation-based group properties are particularly easy to see in Cayley graphs. If two generators commute, the cycle shown in Figure 2.3 will be present at every vertex of the Cayley graph— not surprising, given the earlier discus-

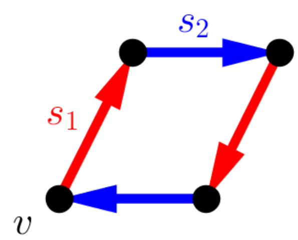

Figure 2.3: If generators s_1 and s_2 commute, this subgraph will appear at every vertex of the Cayley graph.

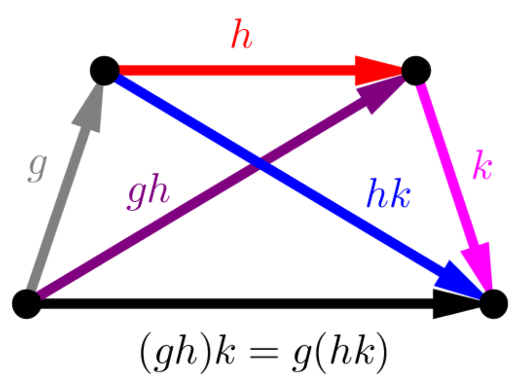

Figure 2.4: A subgraph demonstrating the associative property of group multiplication.

sion of the correspondence between relations and cycles. As a group is commutative if and only if all of its generators commute, in any Cayley graph of a commutative group, one can find the subgraph in Figure 2.3 at every vertex for all $s_1, s_2 \in S$.

If one considers $\vec{\Gamma}_c(G, G)$, group properties that hold for all groups are similarly illustrated by certain subgraphs, as Terrence Tao explores in a blog post [26]. Associativity is one such property. In $\vec{\Gamma}_c(G, G)$, since all group elements are generators, the following fact is true. If you have a path between vertices v and w, then there is an edge from v to w, colored with the word corresponding to the sequence of edge colors in the path. There is also an edge from w to v, colored with the inverse of that word. Given some path on 4 vertices, with edge color sequence ghk, we have the subgraph shown in Figure 2.4, a visual representation that $(gh)k = g(hk)$. In a like manner, we can find a subgraph illustrating the identity $(gh)^{-1} = h^{-1}g^{-1}$ in $\vec{\Gamma}_c(G, G)$. Again

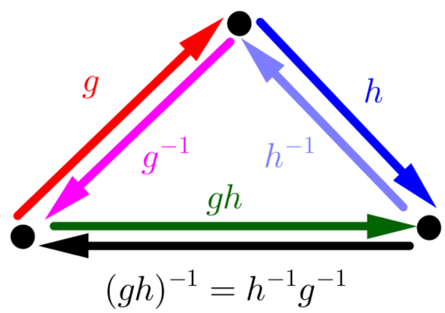

Figure 2.5: A subgraph demonstrating the identity $(gh)^{-1} = h^{-1}g^{-1}$

because all group elements are in the generating set, given an edge (v, w) colored h, there is an edge (w, v) colored h^{-1}. So, if we have a path on three vertices with edge sequence gh, we are guaranteed the existence of the subgraph shown in Figure 2.5, showing that $(gh)^{-1} = h^{-1}g^{-1}$.

2.2 Subgroups and Cosets

Let G be a group, S a generating set of G, and H a subgroup of G. Can H be identified with some structure in $\vec{\Gamma}_c(G, S)$?

The answer to this question is yes, and the structure in $\vec{\Gamma}_c(G, S)$ corresponding to H is a fairly intuitive one. Since H is contained in G as a set, one might expect that H is reflected by some subgraph of $\vec{\Gamma}_c(G, S)$ involving elements of H. Since H is itself a group, it seems reasonable that this subgraph would be a Cayley graph of H with respect to some generating set S'. The edges of this sub-Cayley graph are also edges of $\vec{\Gamma}_c(G, S)$, so the generating set S' of H must be a subset of S. This means that the fact that H is a subgroup of G only "appears" in $\vec{\Gamma}_c(G, S)$ if S contains a generating set of H or, equivalently, $\langle H \cap S \rangle = H$. If this is the case, then for all generating sets S' of H contained in S, $\vec{\Gamma}_c(H, S') \subseteq \vec{\Gamma}_c(G, S)$ (see Figure 2.6).

If $\vec{\Gamma}_c(H, S') \subseteq \vec{\Gamma}_c(G, S)$, we know by vertex transitivity that we can find a subgraph X isomorphic to $\vec{\Gamma}_c(H, S')$ at any vertex of $\vec{\Gamma}_c(G, S)$ (where isomorphism here includes

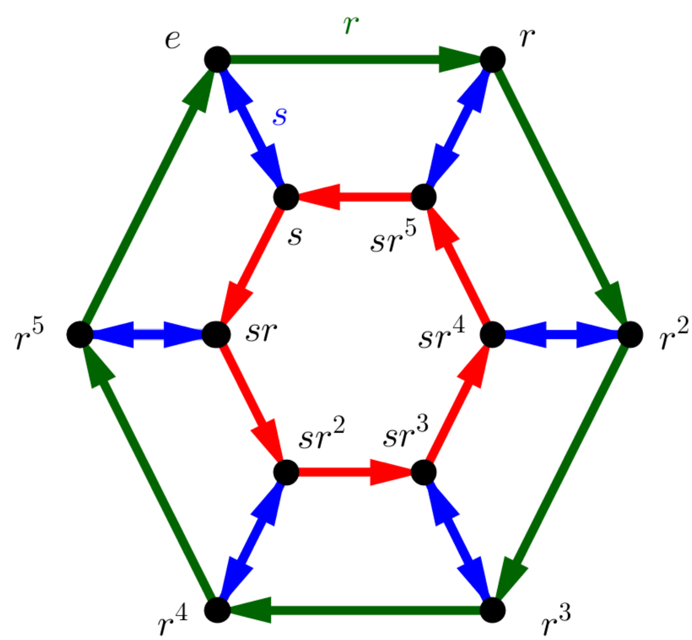

Figure 2.6: Consider two subgroups of D_6 isomorphic to $\mathbb{Z}_6$ and D_3, respectively, $H := \langle r \rangle$ and $F := \langle r^2, s \rangle$. $\vec{\Gamma}_c(H, \{r\})$ is a subgraph of $\vec{\Gamma}_c(D_6, \{r, s\})$, shown here in green. However, $\vec{\Gamma}_c(F, \{r^2, s\})$ is not a subgraph of $\vec{\Gamma}_c(D_6, \{r, s\})$, as $\{r, s\} \cap \{r^2, s\} = \{s\}$ does not generate F.

preservation of edge-colors). If this subgraph X contains a vertex g, then $V(X)$ is made up of all vertices connected to g by a path of S'-colored edges. $E(X)$ consists of all S'-colored edges between the vertices in $V(X)$. Figure 2.7 shows an example. This means that the vertices of X are all of the form $gs'_{\nu_1} \cdots s'_{\nu_m}$, where $s'_{\nu_i} \in S'$ for all $i \in \{1, \ldots, m\}$. Since $S' \subseteq H$, $gs'_{\nu_1} \cdots s'_{\nu_m} = gh$ for some $h \in H$; because S' generates H, gh is a vertex of X for all $h \in H$. This structure may seem familiar to you from group theory: X in fact corresponds to gH, the left coset of H by g. It makes sense from an algebraic standpoint that X is isomorphic to H, since $gH = \lambda_g(H)$ and λ_g is an automorphism of $\vec{\Gamma}_c(G, S)$.

Interpreting left cosets of H as subgraphs of $\vec{\Gamma}_c(G, S)$ isomorphic to $\vec{\Gamma}_c(H, S')$ yields proofs for many basic facts about cosets utilizing combinatorial ideas in addition to algebraic ones. For example, the fact that left cosets of H are either equal or disjoint can be proven using the connectedness of the coset graphs. If the graphs of aH and

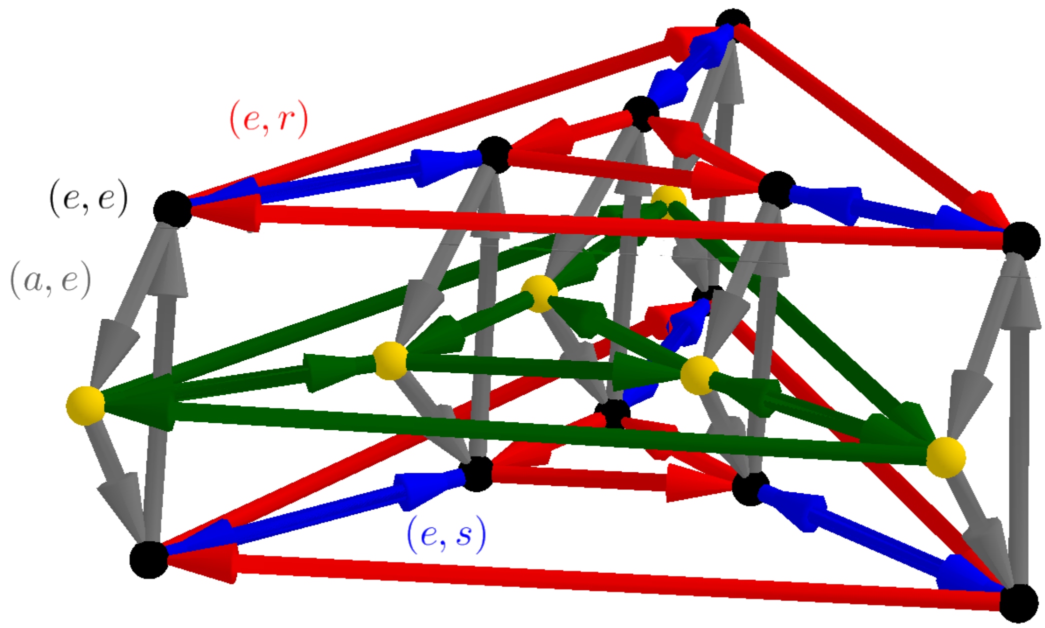

Figure 2.7: Let $G := \langle a | a^3 \rangle$. The graph shown here is $\vec{\Gamma}_c(G \times D_3, \{(a,e),(e,r),(e,s)\})$. Let $S' := \{(e,r),(e,s)\}$ and let $H := \langle S' \rangle \cong D_3$. The subgraph of $\vec{\Gamma}_c(G \times D_3, \{(a,e),(e,r),(e,s)\})$ containing (a,e), the vertices which can be reached from (a,e) by S'-colored edges, and the S'-colored edges among those vertices is shown here by the yellow vertices and green edges. This subgraph corresponds to $(a,e)H$.

bH share a vertex v, then there is a path of S'-colored edges from a to v and from v to any vertex bh in bH. Concatenating these paths gives us a way to write bh as ah', which implies $bh \in aH$. Similarly, we can write any element of aH as an element of bH, so the graphs of both of these cosets, and the cosets themselves, are equal. The characterization of when two cosets are equal can also be proved by thinking of elements of H as paths of S'-colored edges.

The representation of left cosets via subgraphs also gives an illustration of Lagrange's Theorem, or, more generally, the fact that $|G| = |H||G : H|$. The graphs of left cosets of H partition the vertices of $\vec{\Gamma}_c(G,S)$, since each vertex is in exactly one left coset graph. (An example of this partition is illustrated in Figure 2.8.) What we noted earlier about disjoint monochromatic cycles covering the vertices of $\vec{\Gamma}_c(G,S)$ was just a special case of this with $H = \langle s \rangle$. Each coset graph has $|H|$ vertices, being

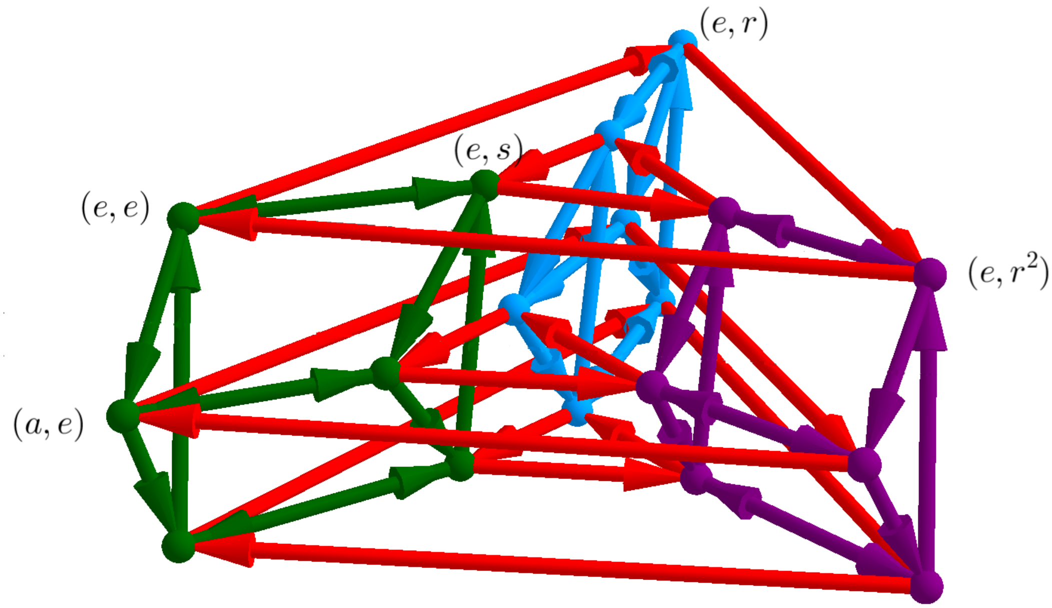

Figure 2.8: Let $G := \langle a | a^3 \rangle$, and $H := \langle (a,e),(e,s) \rangle \leq G \times D_n$; $G \cong \mathbb{Z}_3$ and $H \cong \mathbb{Z}_3 \times \mathbb{Z}_2$. This is $\vec{\Gamma}_c(G \times D_3, \{(a,e),(e,r),(e,s)\})$, with $\vec{\Gamma}_c(H, \{(a,e),(e,s)\})$ shown in green. The blue and purple subgraphs are isomorphic to $\vec{\Gamma}_c(H, \{(a,e),(e,s)\})$, and correspond to left cosets $(e,r)H$ and $(e,r^2)H$ respectively. Note that the subgraphs corresponding to H and its left cosets partition the vertices of $\vec{\Gamma}_c(G \times D_3, \{(a,e),(e,r),(e,s)\})$.

isomorphic to $\vec{\Gamma}_c(H, S')$, so $|G| = |H||G : H|$.

Though the right cosets of H also partition the vertices of $\vec{\Gamma}_c(G, S)$, they do not in general correspond to a particularly meaningful subgraph of $\vec{\Gamma}_c(G, S)$. To find the group elements in Hg, where $g = s_{\nu_1} \cdots s_{\nu_n}$ for $s_{\nu_i} \in S$, one begins at a vertex $h \in H$ and follows the path whose edge color sequence is $s_{\nu_1}, \ldots, s_{\nu_n}$. The vertex at the end of this path is $hg \in Hg$. Usually, the vertices $Hg \subseteq V(\vec{\Gamma}_c(G, S))$ do not form a subgraph isomorphic to $\vec{\Gamma}_c(H, S')$, as shown in Figure 2.9; we really have no reason to expect them to, since right multiplication by a fixed element of G isn't usually an automorphism of $\vec{\Gamma}_c(G, S)$. Indeed, given $h_1 \in H$ and $s' \in S'$, right multiplication by g preserves the edge $(h_1, h_1 s')$ and its color if and only if g is in the centralizer of s'. If we loosen this condition to $(h_1 g, h_1 s' g)$ being an edge colored by some element of S', not necessarily s', the situation is not much improved, as this occurs if and only

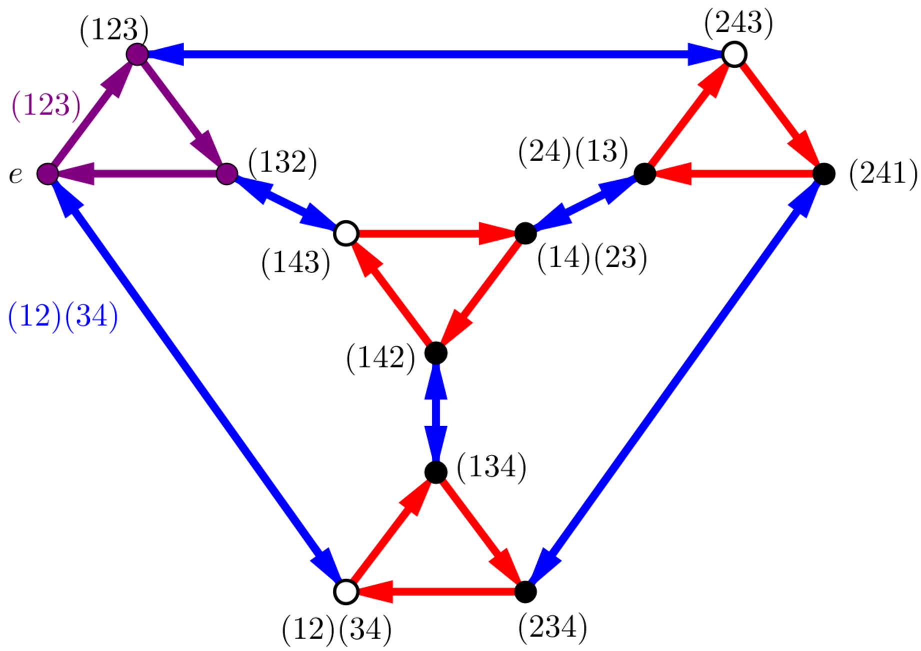

Figure 2.9: Let $H := \langle (123) \rangle \leq A_4$. This digraph is $\vec{\Gamma}_c(A_4, \{(123), (12)(34)\})$, with the vertices and edges of $\vec{\Gamma}_c(H, \{(123)\})$ shown in purple. The hollow vertices make up $H(12)(34)$. As you can see, the only subgraph on these vertices has no edges, so considering these vertices gives us no information about the structure of A_4.

if $g^{-1}sg \in S'$.

2.3 Normal Subgroups and Quotient Groups

Given $\vec{\Gamma}_c(G, S)$ and $H \leq G$ such that $\langle H \cap S \rangle = H$, the above discussion about cosets gives us a way to visually confirm whether or not H is normal in G. The most convenient definition of normality here is $H \trianglelefteq G$ if $gH = Hg$ for all $g \in G$. So $H \trianglelefteq G$ if and only if starting at a vertex $g \in G$ and finding the subgraph consisting of the vertices reachable from g by S'-colored edges and the S'-colored edges between these vertices is exactly the subgraph one arrives at by translating each vertex in $\vec{\Gamma}_c(H, S')$ by the path corresponding to g as a word in S (see Figure 2.10). Of course, equality of gH and Hg doesn't mean that these two processes send an element $h \in H$ to the same vertex; this only occurs if g is in the centralizer of h. This gives us a way to

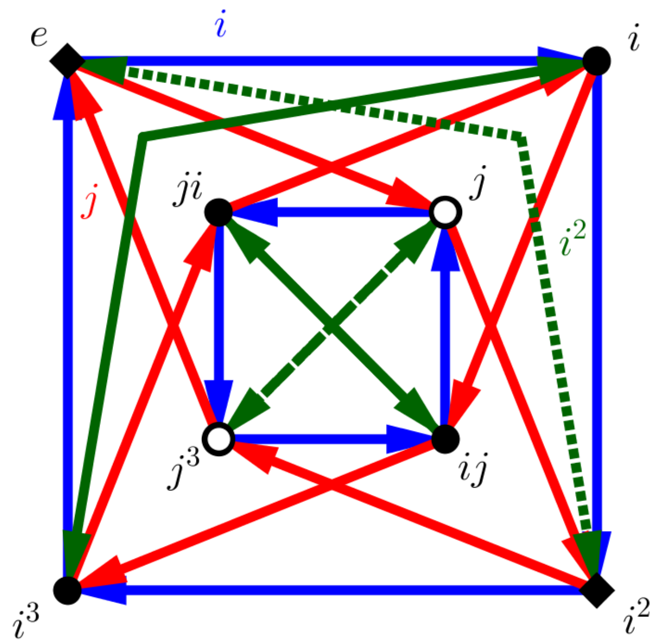

Figure 2.10: The graph above is $\vec{\Gamma}_c(Q, \{i, j, i^2\})$, where Q is the quaternion group. Let $H := \langle i^2 \rangle$. $\vec{\Gamma}_c(H, \{i^2\})$ consists of the diamond-shaped vertices and the dotted edges between them. The subgraph corresponding to jH consists of the two hollow vertices and the dashed edges. If you walk along a j-colored edge from any element of H, you end up in jH, which means that $Hj = jH$. It's not hard to verify that a similar equality holds for any $g \in Q$, implying that H is a normal subgroup.

visually characterize normal subgroups in a Cayley graph. H is normal if and only if for any generator s, s-colored edges originating in the same left coset (which are easily recognizable) all go into a single coset. More specifically, for $s \in S$ and $g \in G$, all s-colored edges starting in gH end in gsH.

If $N \trianglelefteq G$, the quotient group G/N is well-defined; it's fairly natural then, that we turn from normal subgroups and how they appear in Cayley graphs to a discussion of quotient groups. The canonical homomorphism $\alpha : g \mapsto gN$ from G to G/N "describes" how to go from a group to a quotient. α also naturally induces onto maps between the vertex, edge, and edge-color sets of $\vec{\Gamma}_c(G, S)$ and $\vec{\Gamma}_c(G/N, \{sN : s \in S \setminus N\})$, giving us a visual depiction of taking a quotient. These maps, which we also name α in a mild abuse of notation, send vertices g to gN, edges (g, gs) to (gN, gsN), and edge-colors s to sN (where a loop (gN, gN) is identified with the

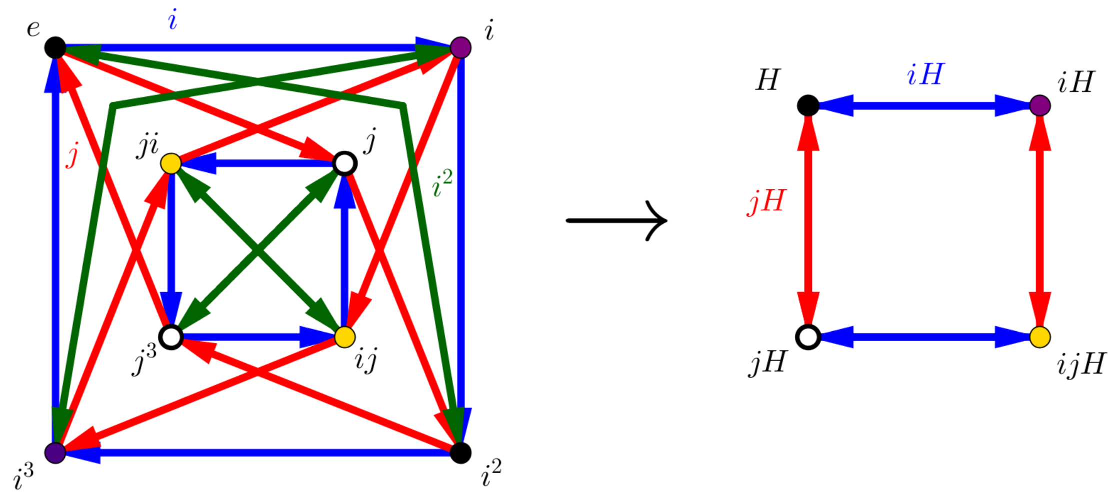

Figure 2.11: "Collapsing" the cosets of $\langle i^2 \rangle$ in $\vec{\Gamma}_c(Q, \{i, j, i^2\})$, the digraph on the right, yields the digraph on the left, which is a Cayley digraph of $Q/\langle i^2 \rangle$. Elements in the same coset of $\langle i^2 \rangle$ are here the same color in $\vec{\Gamma}_c(Q, \{i, j, i^2\})$.

vertex gN). The fact that these maps are well-defined and onto follows immediately from the definition of G/N.

Visually speaking, α sends all vertices in the same coset of N to a single vertex, and makes all directed edges between a pair of cosets into a single edge, recoloring appropriately (see Figure 2.11). We can apply the quotienting process to $\vec{\Gamma}_c(G, S)$ for any normal subgroup of G, but the most structure of the quotient graph is most readily apparent from $\vec{\Gamma}_c(G, S)$ if $\langle S \cap N \rangle = S$. As we explored above, in this case the cosets of N are visually apparent in $\vec{\Gamma}_c(G, S)$ as a collection of disjoint subgraphs isomorphic to $\vec{\Gamma}_c(N, S \cap N)$. α collapses each of these subgraphs to a single vertex. This makes intuitive sense, as the generators that appear as edge-colors in these collapsed subgraphs are exactly those that no longer "matter" in G/N. In taking the quotient, we decide that these generators correspond to the identity element N, so we treat the vertices in $\vec{\Gamma}_c(G, S)$ connected by edges colored by these generators as the same vertex. So, to see a Cayley digraph of G/N within $\vec{\Gamma}_c(G, S)$, one need merely group the vertices according to which coset of N they are in, and then pretend each cluster of vertices is a single vertex, identifying parallel edges.

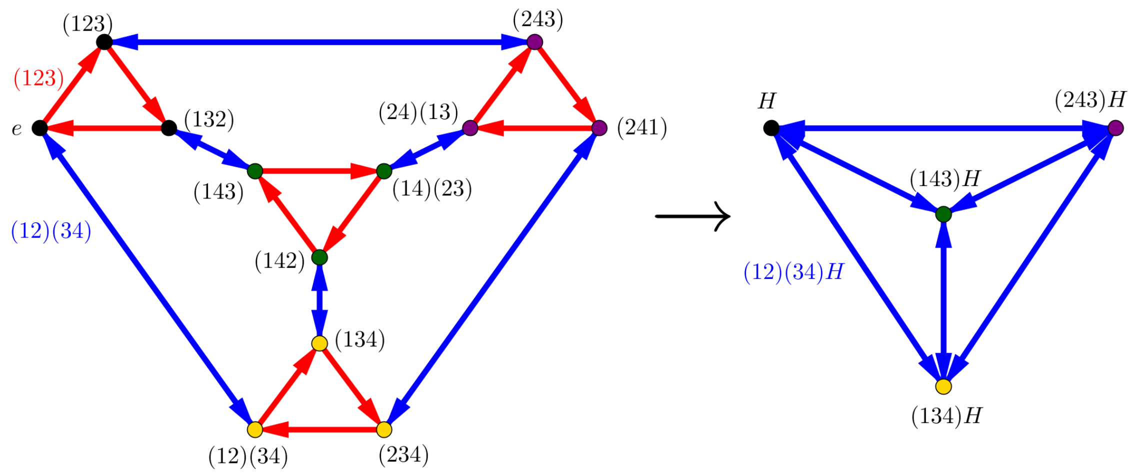

Figure 2.12: If we try to quotient a Cayley digraph by a non-normal subgroup, the resulting digraph isn't a Cayley digraph. Here, if we "quotient" $\vec{\Gamma}_c(A_4, \{(123), (12)(34)\})$, shown on the left, by $H := \langle(123)\rangle$, we get the digraph on the right, which is clearly not a Cayley digraph. Elements in the same left cosets of H are the same color in $\vec{\Gamma}_c(A_4, \{(123), (12)(34)\})$.

The fact that the quotient group G/N is only well-defined for normal subgroups of G is also reflected nicely in Cayley graphs. When N is a normal subgroup of G, we noted earlier that in $\vec{\Gamma}_c(G, S)$, all s-colored edges starting in the same coset end in the same coset. This means that when we collapse each coset of N in $\vec{\Gamma}_c(G, S)$ to a single vertex as described above, no vertex has two edges of the same color coming in to it, or going out. This is a necessary condition for an edge-colored digraph to be a Cayley digraph. But, if we try the same thing with left cosets of a non-normal subgraph H of G, the resulting digraph would not fulfill this condition (see Figure 2.12). Namely, for any generator s not in the normalizer of H, the vertex H would have more than one s-colored edge coming from it in the "quotient" digraph.

2.4 Direct & Semi-direct Products

We turn our attention now to two kinds of group structures involving special normal subgroups, direct and semi-direct products. We can translate the process of taking the product of two groups into an operation on their Cayley digraphs, just as we did for taking the quotient of a group by a normal subgroup. This will allow us to identify when a group is a product of some of its subgroups based on its Cayley digraphs, and to determine which Cayley digraphs display such structure clearly.

First, we need the definition of the Cartesian product of two digraphs.

Definition 2.1. Let $\vec{X} = (V(\vec{X}), E(\vec{X}))$ and $\vec{Y} = (V(\vec{Y}), E(\vec{Y}))$ be digraphs. The *Cartesian product* of $\vec{X}$ and $\vec{Y}$, denoted $\vec{X}\Box\vec{Y}$, is the digraph with vertex set $V(\vec{X}) \times V(\vec{Y})$. For $(x_1, y_1), (x_2, y_2) \in V(\vec{X}\Box\vec{Y})$, there is an edge from (x_1, y_1) to (x_2, y_2) if and only if $x_1 = x_2$ and $(y_1, y_2) \in E(\vec{Y})$ or $y_1 = y_2$ and $(x_1, x_2) \in E(\vec{X})$.

There is a very nice intuitive way to think about the Cartesian product of two digraphs. We can construct $\vec{X}\Box\vec{Y}$ by "inserting" a copy of $\vec{X}$ at every vertex of $\vec{Y}$ (see Figure 2.13), where the vertices planted at $y \in V(\vec{Y})$ are given y as their second coordinate. Then, if $(y_1, y_2) \in E(\vec{Y})$, for all $x \in V(\vec{X})$ we draw an edge from x in the copy of $\vec{X}$ at y_1 to x in the copy of $\vec{X}$ at y_2. Figure 2.13 illustrates this construction. Note that we could also construct $\vec{X}\Box\vec{Y}$ by planting a copy of $\vec{Y}$ at each vertex of $\vec{X}$ and then connecting vertices in different copies of $\vec{Y}$ as described above. This is reflected in the definition of the Cartesian product, since exchanging $\vec{X}$ and $\vec{Y}$ just changes the order of the coordinates; as the edge set of $\vec{X}\Box\vec{Y}$ is symmetric in the coordinates, this doesn't change the structure of the resulting digraph. To put it more formally, $\Box$ is commutative (up to isomorphism).

You may have noticed that taking the Cartesian product of two digraphs seems like an inverse of the quotienting process we described above for Cayley digraphs. Instead of collapsing disjoint isomorphic subgraphs, we're inserting them. Given $\vec{X}\Box\vec{Y}$, we can recover $\vec{Y}$ by collapsing each copy of $\vec{X}$ back into a single vertex and all edges from one copy to another into a single edge; to put it another way, by quotienting with respect to $\vec{X}$. Since the Cartesian product is commutative, we can also recover $\vec{X}$ by quotienting $\vec{X}\Box\vec{Y}$ with respect to $\vec{Y}$. Another important thing to note is that any copy of $\vec{X}$ intersects any copy of $\vec{Y}$ precisely once, since a copy of $\vec{X}$ (resp. $\vec{Y}$)

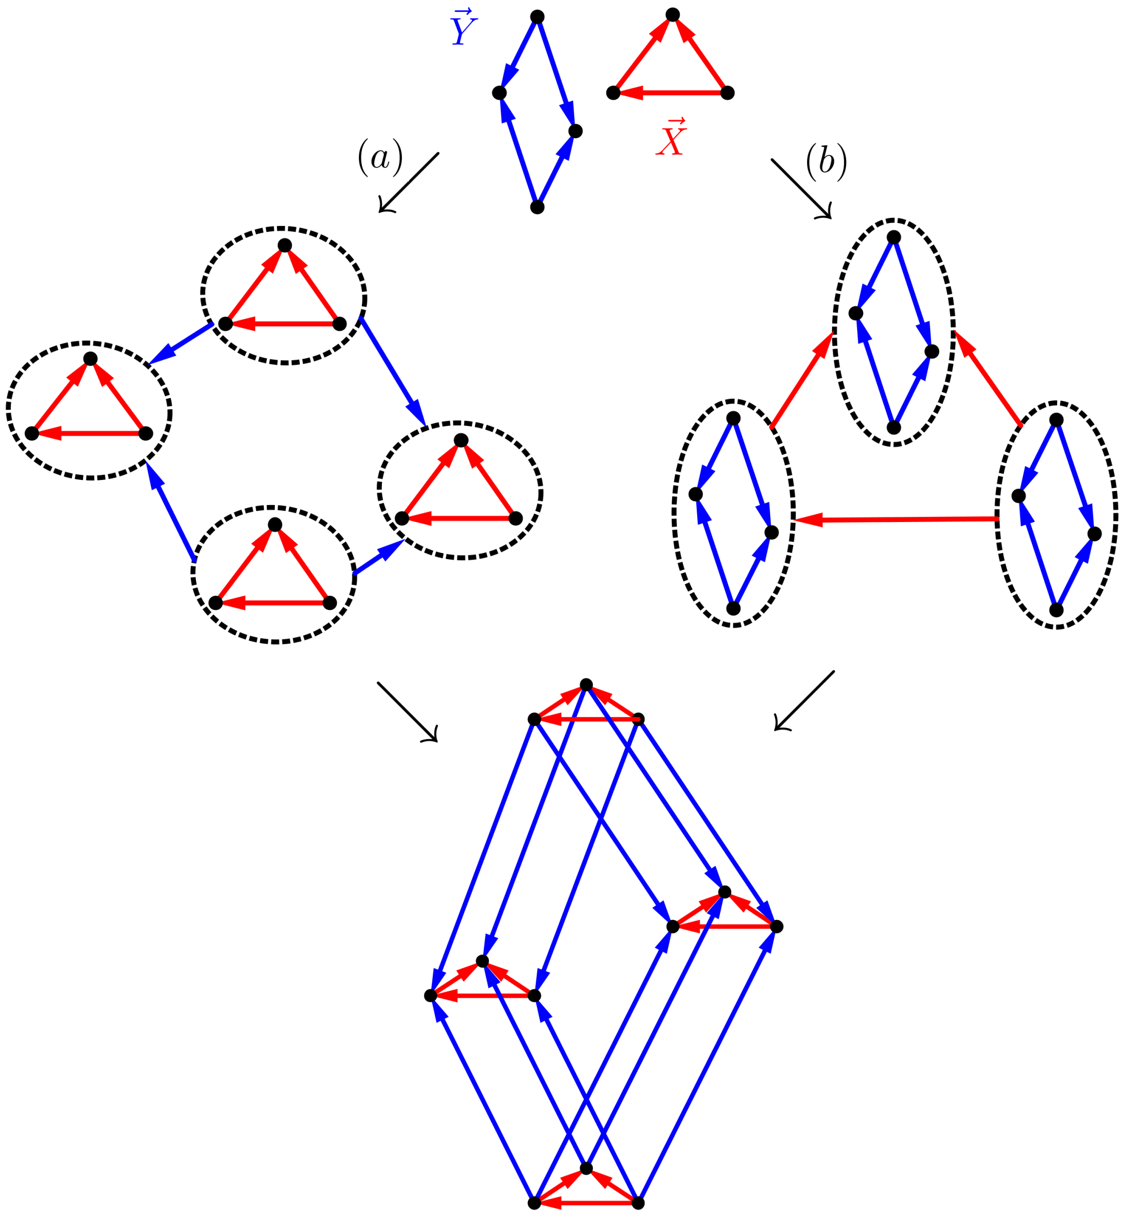

Figure 2.13: Constructing the Cartesian product of a cycle on 3 vertices $\vec{X}$ and a cycle on 4 vertices $\vec{Y}$ in two ways. In method *(a)*, we insert a copy of $\vec{X}$ in every vertex of $\vec{Y}$ and then attaching corresponding vertices in different copies according to the edges of $\vec{Y}$. Method *(b)* is the same as method *(a)*, but with the roles of $\vec{X}$ and $\vec{Y}$ reversed. The result of both methods is $\vec{X}\Box\vec{Y}$, the digraph on the bottom.

consists of all vertices whose second (resp. first) coordinate is some fixed y (resp. x).

If $\vec{X}$ and $\vec{Y}$ are Cayley digraphs, say $\vec{\Gamma}_c(G_1, S_1)$ and $\vec{\Gamma}_c(G_2, S_2)$, the above observations become quite suggestive. (In this case, $\vec{\Gamma}_c(G_1, S_1)\Box\vec{\Gamma}_c(G_2, S_2)$ is edge-colored; an edge inherits its color from the corresponding edge in $\vec{\Gamma}_c(G_1, S_1)$ or $\vec{\Gamma}_c(G_2, S_2)$.) Since quotienting $\vec{\Gamma}_c(G_1, S_1)\Box\vec{\Gamma}_c(G_2, S_2)$ by either $\vec{\Gamma}_c(G_1, S_1)$ or $\vec{\Gamma}_c(G_2, S_2)$ yields a Cayley digraph, we get the feeling that $\vec{\Gamma}_c(G_1, S_1)\Box\vec{\Gamma}_c(G_2, S_2)$ may be the Cayley digraph of some group G in which G_1 and G_2 are normal. Moreover, the fact that the copies of $\vec{\Gamma}_c(G_1, S_1)$ and $\vec{\Gamma}_c(G_2, S_2)$ intersect in only one vertex suggests that the intersection of G_1 and G_2 is trivial. Finally, since every vertex of $\vec{\Gamma}_c(G_1, S_1)\Box\vec{\Gamma}_c(G_2, S_2)$ can be written as an ordered pair of vertices in $\vec{\Gamma}_c(G_1, S_1)$ and $\vec{\Gamma}_c(G_2, S_2)$, we would be able to write every element of G as the product of an element of G_1 and an element of G_2. Having two normal subgroups with trivial intersection whose setwise product is the entire group implies that a group is a direct product, so we suspect that if $\vec{\Gamma}_c(G_1, S_1)\Box\vec{\Gamma}_c(G_2, S_2)$ is a Cayley digraph, its generating group is $G_1 \times G_2$.

This is in fact the case. If G_1 and G_2 are groups with generating sets S_1 and S_2 respectively, and $\iota_j : G_j \to G_1 \times G_2$, $j = 1, 2$, are inclusion maps, then $\vec{\Gamma}_c(G_1 \times G_2, \iota_1(S_1) \cup \iota_2(S_2)) = \vec{\Gamma}_c(G_1, S_1)\Box\vec{\Gamma}_c(G_2, S_2)$. The equality of these two digraphs follows immediately from a comparison of their edge sets, as edges of both are of the form $((g_1, g_2), (g_1s_1, g_2))$ or $((g_1, g_2), (g_1, g_2s_2))$, where $g_j \in G_j$ and $s_j \in S_j$. Edge colors are also the same in each, provided one treats $s_j \in S_j$ and $\iota_j(s_j)$ as the same color. So, in words, the Cartesian product of Cayley digraphs is a Cayley digraph of the direct product of the generating groups. This is illustrated in Figure 2.7, which shows the Cartesian product of $\vec{\Gamma}_c(\mathbb{Z}_3, \{\bar{1}\})$ with $\vec{\Gamma}_c(D_3, \{r, s\})$.

Now that we understand how taking the direct product of two groups translates to the Cayley digraph setting, we can answer the question of which Cayley digraphs of a direct product clearly display the product structure. If $G = G_1 \times G_2$, where $G_1, G_2 \leq G$, then you can "see" that G is a product in $\vec{\Gamma}_c(G, S)$ if $\langle G_j \cap S\rangle =$

G_j for $j = 1, 2$, or, in other words, if your generating set contains a generating set of each of the factors. If this is the case, then $\vec{\Gamma}_c(G, S)$ will contain a Cayley digraph which is the Cartesian product of two Cayley digraphs as a subgraph; namely, $\vec{\Gamma}_c(G_1 \times G_2, (G_1 \cup G_2) \cap S) \leq \vec{\Gamma}_c(G, S)$. Since this subgraph contains all of the vertices of $\vec{\Gamma}_c(G, S)$, its existence indicates that G is a direct product. This answer is not terribly surprising, since the generating set of G containing generating sets of each of its direct factors correspond to the presentations that most emphasize that G is a product. In addition, generating sets of this type are required for one to clearly see in $\vec{\Gamma}_c(G, S)$ that G_1 and G_2 are (normal) subgroups of G. It would be difficult to deduce from $\vec{\Gamma}_c(G, S)$ that G was a direct product if this fact were not visible, so it's certainly reasonable that $\langle G_j \cap S \rangle = G_j$ for $j = 1, 2$ is necessary for the direct product structure of G to be apparent in $\vec{\Gamma}_c(G, S)$.

Now, to semidirect products. Just as with direct products, we first examine external semi-direct products and how they can be viewed as a combination of Cayley digraphs, and then move to recognizing internal semi-direct products via Cayley digraphs.

Let H and N be groups and $\theta : H \to Aut(N)$ a homomorphism. For convenience, we write $\theta(h_2)$ as θ_{h_2}. Recall that the semi-direct product of N by H (with respect to θ), denoted $H \ltimes_\theta N$, is a group whose set of elements is $H \times N$ and whose operation is $(h_1, n_1)(h_2, n_2) = (h_1h_2, \theta_{h_2}(n_1)n_2)$ (see [11, Sec. 5.5]). Semi-direct products are generalizations of direct products, as the direct product of H and N is $H \ltimes_\beta N$ where β is the trivial homomorphism. One of the ways in which the two products differ is that only N is necessarily a normal subgroup $H \ltimes_\theta N$; H is always a subgroup $H \ltimes_\theta N$ but frequently is not normal.

If we are given two groups H and N, generating sets S_H and S_N of each (in other words, $\vec{\Gamma}_c(H, S_H)$, $\vec{\Gamma}_c(N, S_N)$), and some homomorphism $\theta : H \to Aut(N)$, how can we construct a Cayley digraph of $H \ltimes_\theta N$? First, we need to choose a generating set.

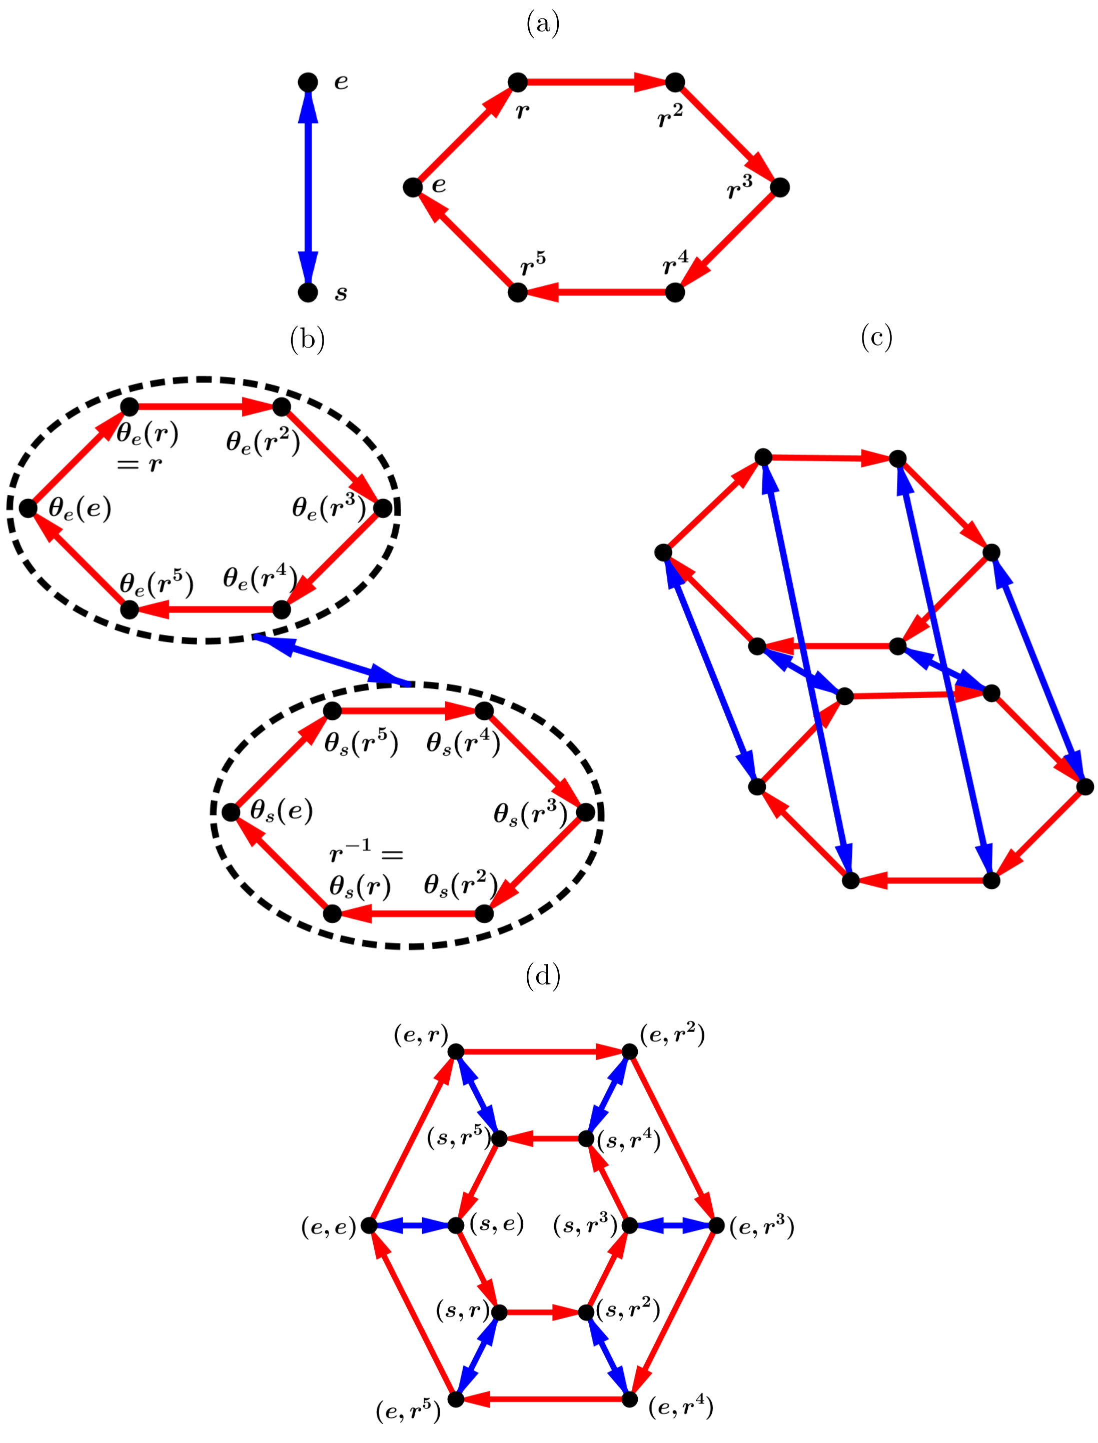

Figure 2.14: **(a)** Let $G := \langle s | s^2 \rangle$, $F := \langle r | r^6 \rangle$ and $\theta : G \to Aut(F)$ be defined by $\theta_e(r) = r$ and $\theta_s(r) = r^{-1}$. $\vec{\Gamma}_c(G, \{s\})$ and $\vec{\Gamma}_c(F, \{r\})$ are shown here. **(b)** We place a copy of $\vec{\Gamma}_c(F, \{r\})$ at each vertex g of $\vec{\Gamma}_c(G, \{s\})$. **(c)** Then, we attach a vertex in a copy of $\vec{\Gamma}_c(F, \{r\})$ to its images under θ_g in other copies, according to the edges of $\vec{\Gamma}_c(G, \{s\})$. **(d)** The resulting digraph is $\vec{\Gamma}_c(G \ltimes_\theta F, \{(s, e), (e, r)\})$, which is isomorphic to $\vec{\Gamma}_c(D_6, \{r, s\})$.

Given the information we have, the natural choice here is $S = \iota_H(S_H) \cup \iota_N(S_N)$, where ι_H and ι_N are inclusion maps from H and N, respectively, to $H \ltimes_\theta N$. To understand the structure of $\vec{\Gamma}_c(H \ltimes_\theta N, S)$, we first note that it contains a number of disjoint subgraphs isomorphic to $\vec{\Gamma}_c(N, S_N)$, comprised of all vertices whose first coordinate is some fixed $h \in H$. As with direct products, then, to create $\vec{\Gamma}_c(H \ltimes_\theta N, S)$, we should insert a copy of $\vec{\Gamma}_c(N, S_N)$ at each vertex of $\vec{\Gamma}_c(H, S_H)$. (Since $N \trianglelefteq H \ltimes_\theta N$ and quotienting by N leaves H, this is not too surprising.) We still use the edges of $\vec{\Gamma}_c(H, S_H)$ as a guide for edges between different copies of $\vec{\Gamma}_c(N, S_N)$, in the sense that there is an edge from a vertex n_1 in the copy of $\vec{\Gamma}_c(N, S_N)$ at h_1 and a vertex n_2 in the copy of $\vec{\Gamma}_c(N, S_N)$ at h_2 only if (h_1, h_2) is an edge of $\vec{\Gamma}_c(H, S_H)$. But here, instead of connecting n_1 in the copy of $\vec{\Gamma}_c(N, S_N)$ at h_1 to itself in the copy at h_2, we connect n_1 in the copy at h_1 to $\theta_{h_2}(n_1)$ in the copy at h_2. This edge is then colored the same color as (h_1, h_2) in $\vec{\Gamma}_c(H, S_H)$. Phrased differently, the edges of $\vec{\Gamma}_c(H, S_H)$ tell us which cosets of N should have edges between them, and what their direction and color should be; for two cosets that should have edges between them, θ tells us exactly which vertices in one the vertices of the other should be connected to. See Figure 2.14 for an illustration of this process.

We should note the properties of $\vec{\Gamma}_c(H \ltimes_\theta N, S)$ that reflect what we know of the structure of semi-direct products. As H is a subgroup, the Cayley digraph contains $|N|$ disjoint subgraphs isomorphic to $\vec{\Gamma}_c(H, S_H)$, corresponding to left cosets of $\iota_H(H)$. These subgraphs are not as immediately visible in $\vec{\Gamma}_c(H \ltimes_\theta N, S)$ as those corresponding to cosets of N, since they aren't comprised of vertices with a fixed second coordinate and all possible first coordinates. Instead, the second coordinates vary depending on how the first coordinates are related. However, using the strategy outlined in the section on subgroups of choosing a starting vertex and then collecting all the vertices and edges in S_H-colored paths starting at that vertex, one can still fairly easily find these subgraphs. Lastly, the copies of $\vec{\Gamma}_c(N, S_N)$ and $\vec{\Gamma}_c(H, S_H)$

intersect in only one vertex, reflecting that $\iota_N(N) \cap \iota_H(H) = \emptyset$.

We turn now to the question of which Cayley digraphs of a group G would make visually apparent its semi-direct product structure. The answer is quite similar to the direct product case: exactly those Cayley digraphs that show that its (semi-direct) factors H and N are subgroups of G. In other words, those Cayley digraphs whose generating set S contains generating sets of each of the factors. In this case, once the two factors H and N have been identified and the vertices of $\vec{\Gamma}_c(G, S)$ relabelled as elements of the internal semi-direct product, $\theta : H \to Aut(N)$ can be recovered from the Cayley digraph. For any $(h, e) \in H \cap S$, to find θ_h, just look at the (h, e)-colored edges coming from vertices of the type (e, n) for $n \in N$. These edges will go from (e, n) to $(e, \theta_h(n))$, completely describing θ_h. As θ is a homomorphism and $H \cap S$ generates H, this is enough to determine θ completely.

Chapter 3

Group Representations

Cayley (di)graphs give us a way to associate a graph to a group. This association is relatively strong: as we saw in the previous chapters, there are many easily discernible connections between an edge-colored Cayley digraph and its generating group. Not the least of these is the fact that the edge-colored Cayley digraph's automorphism group is isomorphic to its generating group (see Proposition 1.2). When we pass to Cayley digraphs by dropping edge colors or to Cayley graphs by dropping edge colors and directions, we lose some of the clear visual correspondences between (di)graph structure and group structure, but we retain a relationship between the generating group and the (di)graph's automorphism group. Namely, the former appears as a regular subgroup of the latter.

But there are other, more natural ways to connect groups and (di)graphs. After all, the automorphisms of a (di)graph form a group, so one way to do this is to associate a (di)graph with its automorphism group. Indeed, if G is a group and X is a (di)graph, we call X a *(di)graph representation* of G if $A(X) \cong G$. In this chapter, we'll discuss some of the aspects of (di)graph representations of groups and the way they relate to Cayley (di)graphs.

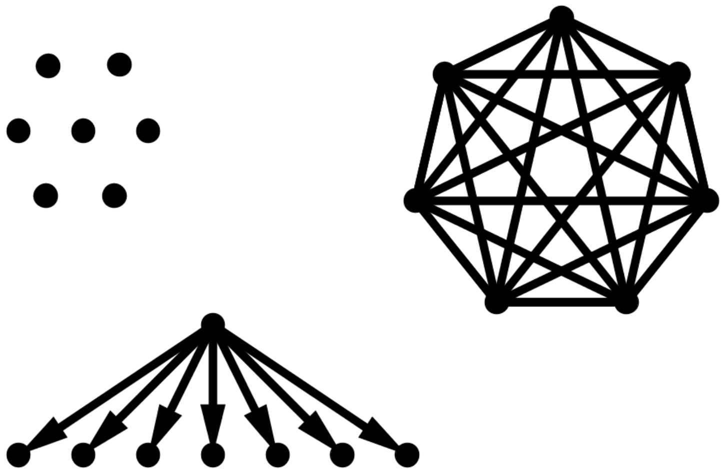

Figure 3.1: Some (di)graph representations of S_7.

3.1 (Di)graph Representations and the Representation Problem

Just as a group has many different Cayley (di)graphs, different (di)graphs can be representations of the same group. S_n, for example, is represented by the complete graph on n vertices, the graph consisting of n vertices and no edges, and the directed star on $n+1$ vertices (see Figure 3.1). We've seen that the Cayley (di)graphs of a given group can be very structurally different; the same is true of (di)graph representations of a given group, as illustrated by the example of S_n. However, isomorphic (di)graphs are always representations of the same group, as they have isomorphic automorphism groups, and a (di)graph is a representation of only one group (up to isomorphism). This is quite different than the situation with Cayley (di)graphs, as two isomorphic Cayley (di)graphs could have nonisomorphic generating groups.

The Cayley (di)graphs and (di)graph representations of groups differ from each other in a number of key ways. For Cayley (di)graphs, we start with a group and get a (di)graph by the process described in Chapter 1; for graph representations, we

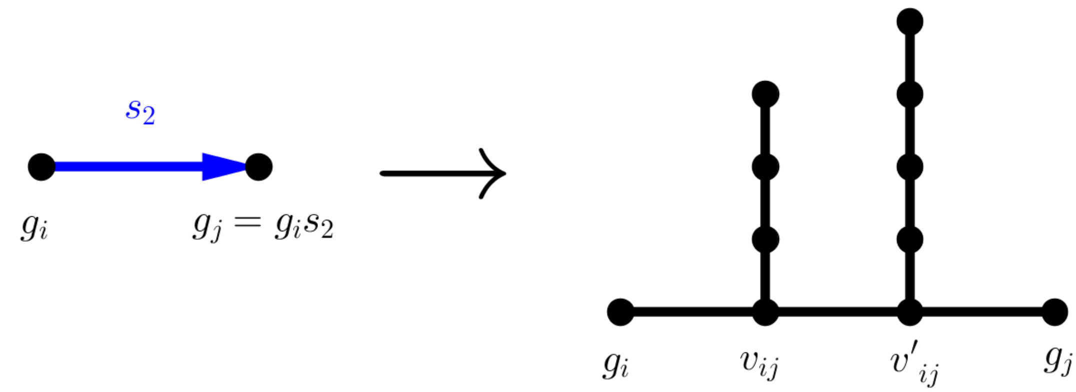

Figure 3.2: In Theorem 3.1, an edge $(g_i, g_i s_2)$ in X', shown on the left, is replaced with the subgraph on the right in X.

instead start with a (di)graph and “get” a group by considering its automorphisms. So though it’s obvious that every group has a Cayley (di)graph, it is not necessarily the case that every group has a (di)graph representation. In other words, it’s possible that some group can’t be realized as the symmetry group of a (di)graph.

The question of which (finite) groups can be represented by graphs was first stated by König in 1936 [19]. Two years later, Frucht [12] gave a very nice answer: every finite group. His proof, which we give below, is a nice application of Cayley digraphs to the representation problem. The idea is to take an edge-colored Cayley digraph of a given group, and then replace edge-colors and directions with small asymmetric graphs to turn the Cayley digraph into a graph without changing its automorphism group.

Theorem 3.1 ([12])**.** *Let G be a finite group. There exists a (finite) graph X such that $A(X) \cong G$.*

Proof. ([7])

If $G = 1$, then $X = \{\{v_1\}, \{\}\}$ is a graph with $A(X) = G$.

Now suppose $G = \{g_1, \ldots, g_n\}$, $n \geq 2$ and let $S = \{s_1, \ldots, s_t\}$, $1 \leq t \leq n$, be some generating set of G. Consider $X' := \vec{\Gamma}_c(G, S)$. By Proposition 1.2, $A(X') \cong G$.

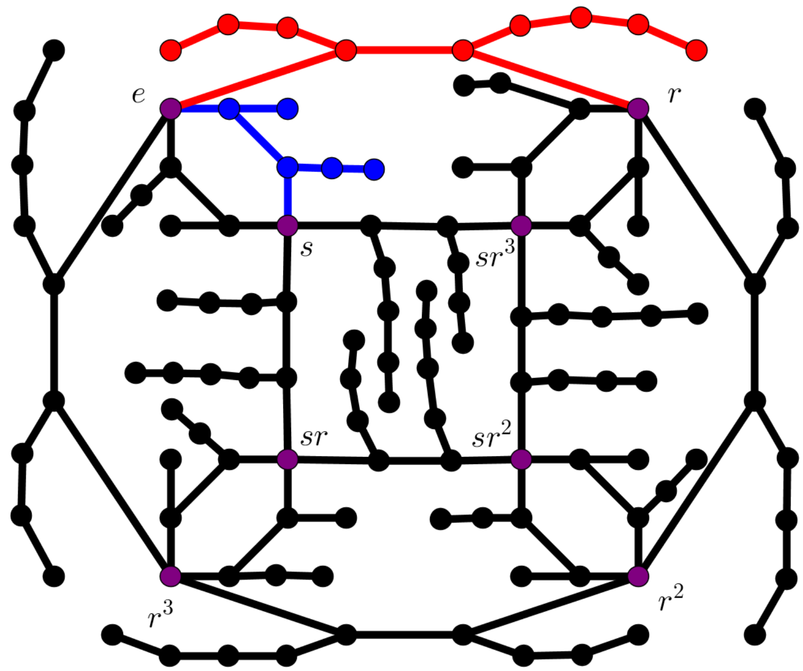

Figure 3.3: A graph representation of D_4, obtained by the process outlined in Theorem 3.1 with $X' = \vec{\Gamma}_c(D_4, \{s, r\})$, and $s_1 := s$ and $s_2 := r$. The subgraphs corresponding to s- and r-colored edges are shown in blue and red, respectively. The vertices of X' are shown in purple.

We create X in the following manner. Suppose $(g_i, g_i s_k)$ is an edge of X', where $g_i s_k = g_j$. We add two vertices v_{ij} and v'_{ij} to the vertex set. Then we delete this s_k-colored directed edge and add the undirected path $g_i, v_{ij}, v'_{ij}, g_j$ to the edge set. Finally, we insert an undirected path on $2k$ vertices beginning at v_{ij} and an undirected path on $2k+1$ vertices beginning at v'_{ij} (adding $2k-1$ vertices to the vertex set for the first path and $2k$ for the second). Note that we do not connect the vertices in these paths to any other vertices in the graph. See Figure 3.2 for an illustration of this process. We do this for every edge of X' and define the resulting graph to be X (see Figure 3.3).

Note that X contains the edge information of X'. If g_i and g_j are adjacent in X', they're connected by a path on 4 vertices in X, say g_i, u, v, g_j, where u and v are the starting vertices (or "roots") of two paths with no edges to the rest of the graph. Let $Y \subseteq X$ consist of the path g_i, u, v, g_j and the paths rooted at u and v. The length of the rooted paths tells us the direction of the g_i, g_j edge in X', as the root of the shorter path is adjacent to the tail of the g_i, g_j edge. The number of vertices in the even path (divided by two) gives us the color of the edge.

It is not hard to see that an automorphism of X' extends to an automorphism of X, since two edges of the same color in X' correspond to two identical subgraphs in X. Each automorphism of X induces also an automorphism of X'. First, note that if $\alpha \in A(X)$, then $\alpha(V(X')) = V(X')$, since the other vertices of X are in a path with no edges to the rest of the graph. Suppose g_i and g_j are adjacent in X'. α must take u and v (from the above paragraph) to two vertices that are the roots of two paths, where the path rooted at $\alpha(u)$ (resp. $\alpha(v)$) is the same length as that rooted at u (resp. v). This means exactly that $\alpha(g_i)$ and $\alpha(g_j)$ are adjacent in X' and the edge connecting them is of the same color and direction as that connecting g_i and g_j. If g_i and g_j aren't adjacent in X', they aren't in a subgraph isomorphic to Y in X. This implies their images under α also won't be in a subgraph isomorphic to Y, so $\alpha(g_i)$ and $\alpha(g_j)$ aren't adjacent in X'. In other words, the function induced by α on $V(X')$ is a bijective function respecting adjacency and non-adjacency.

$\square$

Since graphs can be interpreted as digraphs with symmetric edge sets, Frucht's result has the following corollary.

Corollary 3.1. *Let G be a finite group. There exists a (finite) digraph $\vec{X}$ such that $A(\vec{X}) \cong G$.*

The mathematical impact of König's question and Frucht's answer is worth briefly discussing, though it is not strictly related to our interest in graph representations. Frucht's graph representation result was the beginning of a larger avenue of investigation known as the *representation problem.* The representation problem can be stated as follows: given some class C of mathematical objects (graphs, fields, etc.), which groups appear as automorphism groups of elements of C? A class is called *universal* if the answer to this question is "all groups"; it is called *finitely universal* if all finite groups can be realized as automorphisms groups of finite objects in the class. In this language, we can restate Theorem 3.1 as "the class of all graphs is finitely universal." Babai gives an overview of the representation problem and its development [3, Sec. 4]. Roughly speaking, most "reasonably broad" classes of combinatorial or algebraic objects are universal, where "reasonably broad" classes include commutative rings, topological spaces, and Steiner triple systems. The finite universality of graphs is an important element of many universality proofs; as Babai mentions, a typical universality proof involves somehow encoding a graph with a given automorphism group into another mathematical object.

Once Frucht showed that (di)graphs are finitely universal, attention turned to examining the existence and characteristics of (di)graph representations of groups satisfying additional combinatorial or algebraic constraints. For example, Frucht [13] showed that every finite group has a (finite) graph representation where every vertex has precisely 3 neighbors. This direction of inquiry was motivated by a number of interesting combinatorial results involving these graphs, which are called *cubic*. The fact that the (quite restricted) class of cubic graphs is also finitely universal is perhaps quite surprising; the proof is not as intuitively clear as the proof of Theorem 3.1 and does not involve Cayley graphs, so we omit it here.

Also of interest were "very symmetric" (di)graph representations of groups. By this, we mean representations with the minimum possible number of vertex orbits under the action of $A(X)$. Since the orbit of a vertex v with respect to the action of $A(X)$ consists essentially of those vertices which are "indistinguishable" from v, fewer orbits implies more symmetry. Most groups have quite symmetric graph representations: Babai [1] provides a construction for a representation with two orbits for finite non-cyclic groups, and Sabidussi [25] gives a representation with two orbits for cyclic groups of prime power order larger than 6. The groups that have a 1-orbit (i.e. vertex-transitive) (di)graph representation were found largely as a corollary of a representation problem with additional constraints, our next topic of discussion.

3.2 Graph and Digraph Regular Representations

We now turn to the intersection of Cayley (di)graphs and (di)graph representations. Given that both kinds of (di)graphs are a way to "translate" a group into the setting of (di)graphs, it seems reasonable to ask when the two notions coincide. In other words, when is a (di)graph both a Cayley (di)graph and a (di)graph representation of a given group G?

This questions has a very straightforward answer. Sabidussi's theorem (Proposition 1.3) tells us that X is a Cayley (di)graph of a group G if and only if $\widetilde{G}$, the left regular representation of G, is a regular subgroup of $A(X)$. Recall that a subgroup of $A(X)$ is regular if it acts transitively on $V(X)$ and all stabilizers are trivial. If X is also a representation of G, then $A(X) \cong G$. So, X is both a Cayley (di)graph and a representation of a group G if and only if its automorphism group is regular, in which case $\widetilde{G} = A(X)$. We call such (di)graphs *digraph regular representations* (DRRs) or *graph regular representations* (GRRs).

A (di)graph regular representation X of a group G is something of an extremal element of both the set of Cayley (di)graphs of G and the set of (di)graph representations of G. X is a "maximally symmetric" (di)graph representation of G, because it's vertex-transitive. It is the largest vertex-transitive representation of G, as the Orbit-Stabilizer Lemma tells us that the size of a vertex-transitive (di)graph Y is a divisor of $|A(Y)|$. In addition, X is the vertex-transitive representation on which the automorphism group acts most simply, since the automorphism group acting regularly on the vertex set means that for any two vertices v, w, there is a unique automorphism

α such that $\alpha(v) = w$. X is also a minimally symmetric Cayley (di)graph, in the sense that it has as few automorphisms as Cayley (di)graph of G can possibly have. This means that no automorphisms were gained when edge-colors or edge-colors and directions were removed from the edge-colored Cayley digraph with the same generating set. In some sense, then, we haven't lost any information about the group in the removal process. So we can think of (di)graph regular representations as particularly "good" reflections of their generating groups, as compared to other Cayley (di)graphs.

Our next question is more difficult. Which (finite) groups have a DRR? Or a GRR? Inquiry into the second question began in the 1960s and continued until Godsil [15] determined in 1981 that all (finite) non-solvable groups have a GRR, the status of all other groups being by that time known. The DRR question was easier to answer, which seems intuitively reasonable given that removing edge-colors from an edge-colored Cayley digraph generally introduces fewer automorphisms than removing both edge-colors and directions. Babai [2] determined the groups with a DRR in 1980, using results emerging from work done on the GRR problem. Because of this, we start with GRRs, though the GRR result is more complicated. The intent of our discussion of the two problems is to give the reader an idea of the general development of the problem, rather than to detail how certain groups were proved to have (or not have) GRRs. As a result, we focus largely on the results that motivate these proofs and generally omit the proofs themselves.

3.2.1 Groups with GRRs

We begin with the complete list of groups without a GRR, as given by Godsil in the paper containing the final proofs of the GRR problem [15]. To give this list concisely, we mention two definitions here. The *exponent* of a finite group G is the least integer n such that for all $g \in G$, $g^n = e$ [11, pg. 165]. A *generalized dicyclic group* is a group generated by an abelian group A of exponent greater than two and an element $b \notin A$, where $b^4 = e$, $b^2 \in A \setminus \{e\}$, and $b^{-1}ab = a^{-1}$ for all $a \in A$.

Theorem 3.2. *A finite group G does not admit a graph regular representation if and only if it is:*

i.) *An abelian group of exponent greater than two [6];*

ii.) *A generalized dicyclic group. [20, 29];*

iii.) $\mathbb{Z}_2^k$, *where* $k \in \{2, 3, 4\}$ *[17];*

iv.) D_k *where* $k \in \{3, 4, 5\}$ *[29];*

v.) A_4 *[28];*

vi.) $\langle a, b, c | a^2, b^2, c^2, abc(bca)^{-1}, abc(cab)^{-1} \rangle$ *[28];*

vii.) $\langle a, b | a^8, b^2, b^{-1}aba^{-5} \rangle$ *[21];*

viii.) $\langle a, b, c | a^3, b^3, c^2, a^{-1}b^{-1}ab, (ac)^2, (bc)^2 \rangle$ *[28];*

ix.) $\langle a, b, c | a^3, b^3, c^3, a^{-1}c^{-1}ac, b^{-1}c^{-1}bc, c^{-1}a^{-1}b^{-1}ab \rangle$ *[22];*

x.) $Q \times \mathbb{Z}_k$ *where* $k \in \{3, 4\}$ *[27].*

The references given above are for sufficiency; the necessary direction was proven for solvable groups by Hetzel [16] and for non-solvable groups by Godsil [15].

There is another, more conceptually transparent way to phrase the GRR result, which requires some additional definitions. Because we are dealing with Cayley graphs, from now on we assume every generating set S is *symmetric*, meaning that $S^{-1} = S$. No generality is lost in this assumption, as every edge of the form $\{g, gs^{-1}\}$ can be written as $\{gs^{-1}, (gs^{-1})s\}$.

Definition 3.1. [29] Let G be a finite group. G is said to be in *Class I* if G has a GRR. G is said to be in *Class II* if for each generating set S of G, there exists a non-identity group automorphism $\phi \in Aut(G)$ such that $\phi(S) = S$.

Watkins [29] notes that Class I and Class II are disjoint. Indeed, suppose S is a generating set of a group G and let $\Gamma := \Gamma(G, S)$. Suppose we have a non-identity group automorphism ϕ which fixes S setwise. ϕ is also a non-identity permutation of $V(\Gamma)$ which fixes the identity element. ϕ is in fact an automorphism of Γ, as the edge $\{g, gs\}$ is mapped to $\{\phi(g), \phi(gs)\} = \{\phi(g), \phi(g)\phi(s)\}$. $\phi(s) \in S$ by assumption, so $\{\phi(g), \phi(g)\phi(s)\} \in E(\Gamma)$. ϕ^{-1} also clearly preserves adjacency, which implies $\phi \in A(\Gamma)$. In particular, $\phi \in A(X)_e$, the stabilizer of the identity element. (It is not hard to show that the group automorphisms fixing S setwise in fact form a subgroup of this stabilizer). This implies Γ is not a GRR of G, or even a representation of G, as $|A(\Gamma)| = |Orb_{A(\Gamma)}(e)| \cdot |A(\Gamma)_e| \geq |V(\Gamma)| \cdot 2 > |G|$.

The fact that the two classes are disjoint gives us a sufficient condition for a group not having a GRR, namely the existence of group automorphisms fixing each generating set setwise (the automorphism may depend on the generating set). In the same paper, Watkins conjectured that this condition is also necessary, or, in other words, that the union of Class I and Class II contains all finite groups. His conjecture was proved to be true by the characterization of groups with GRRs given above. Thus, we have another way to think of the GRR result, which gives more insight on the structure of groups without GRRs.

Theorem 3.2 (Version 2)**.** *A finite group G does not admit a graph regular representation if and only if for each symmetric generating set S of G, there exists a non-identity group automorphism $\phi \in Aut(G)$ such that $\phi(S) = S$*

This characterization is particularly interesting in the context of the broader aim of this thesis, to investigate the connections between groups and their Cayley (di)graphs.

As we noted above, a group automorphism ϕ fixing a generating set S of a group G is in fact a graph automorphism of $\Gamma(G, S)$. Such group automorphisms are the only ones which can be graph automorphisms, as all group automorphisms fix the identity element. The neighbors of e in $\Gamma(G, S)$ are precisely the elements of S. If $\phi \in Aut(G)$ doesn't fix S setwise, ϕ doesn't preserve the edges involving e and thus is not a graph automorphism of $\Gamma(G, S)$. The second version of Theorem 3.2 tells us that these automorphisms, the ones which are automorphisms of both a group and one of its Cayley graphs, are in some sense the crux of the GRR question. The groups which do not have GRRs are precisely those groups that share a nontrivial automorphism with each of their Cayley graphs. A note of caution about the converse of this statement: while it is true that a group with a GRR has Cayley graphs with which it does not share nontrivial automorphisms and its GRRs are among these Cayley graphs, it is not the necessarily case that every Cayley graph that has no nontrivial automorphism in common with its generating group is a GRR.

We turn now to the strategies used to approach the GRR problem. Watkins' two classes suggest a method to show a group doesn't have a GRR: show it is in Class II. Indeed, this is largely the technique used for the groups without a GRR listed in Theorem 3.2, though Nowitz [20] gives a more direct proof that the generalized dicyclic groups have no GRRs. This method is sometimes quite easy, as is the case for abelian groups of exponent greater than 2. The map which takes an element to its inverse is a nontrivial group automorphism which fixes every symmetric generating set. The generalized dihedral groups have a similarly general nontrivial automorphism which fixes all symmetric generating sets, given in [29]. More frequently, however, the desired group automorphism depends on the generating set chosen.

Similarly, there is one method used overwhelmingly to show that a certain Cayley graph $\Gamma(G, S)$ *is* a GRR of its generating group. It is motivated by the following result of Watkins and Nowitz.

Proposition 3.1 ([21])**.** *Let G be a finite group, S a symmetric generating set of G and $R \subseteq S$. Let $\Gamma := \Gamma(G, S)$. Suppose that $A(\Gamma)_e = A(\Gamma)_r$ for all $r \in R$. Then for all $\tilde{r} \in \langle R \rangle$, $A(\Gamma)_e = A(\Gamma)_{\tilde{r}}$.*

Or, in words: if all graph automorphisms fixing the identity vertex fix R pointwise, then they also fix the group generated by R pointwise.

Proof. Note that by the transitivity of the action of $A(\Gamma)$ on $V(\Gamma)$, all vertex stabilizers are conjugate subgroups and thus the same size. Hence, it suffices to show inclusion.

Let $g \in G$ and suppose $\alpha \in A(\Gamma)$ is in the stabilizer of g. Recall that $\lambda_g \in A(\Gamma)$ is the automorphism that maps h to gh. $\lambda_g^{-1}\alpha\lambda_g \in A(\Gamma)_e$, as $\lambda_g^{-1}\alpha\lambda_g(e) = g^{-1}\alpha(g) = e$. By hypothesis, this implies $\lambda_g^{-1}\alpha\lambda_g(r) = r$ for all $r \in R$. In other words, $\alpha(gr) = gr$.

Suppose $\beta \in A(\Gamma)_e$, and let $\tilde{r} = r_{b_1} \cdots r_{b_n} \in \langle R \rangle$. By hypothesis, $\beta(r_{b_1}) = r_{b_1}$. By the above observation, $\beta(r_{b_1} \cdots r_{b_i}) = r_{b_1} \cdots r_{b_i}$ implies $\beta(r_{b_1} \cdots r_{b_{i+1}}) = r_{b_1} \cdots r_{b_{i+1}}$. Applying this for $i = 1, \ldots, n-1$ gives that $\beta(\tilde{r}) = \tilde{r}$.

$\square$

Another way to think of this proof is that $\alpha \in A(\Gamma)_g$ fixes those neighbors h of g such that (g, h) is an r-colored edge in $\vec{\Gamma}_c(G, S)$. Thus $\alpha \in A(\Gamma)_h$, and we conclude that α fixes the neighbors k of h such that (h, k) is an r-colored edge in $\vec{\Gamma}_c(G, S)$. Continuing this process, we get that α fixes all vertices reachable from g in $\vec{\Gamma}_c(G, S)$ by a path in r-colored edges. From Section 2.2, we know that these vertices are precisely $g\langle R \rangle$. When $g = e$, the fixed vertices are just $\langle R \rangle$.

How does Proposition 3.1 help us find GRRs? If R is a subset of a generating set S of G satisfying the conditions of the proposition and in addition $\langle R \rangle = G$, then Proposition 3.1 implies that $\Gamma(G, S)$ is a GRR of G. Any graph automorphism fixing e fixes all other vertices and thus $A(\Gamma(G, S))_e$ is trivial. The vertex stabilizers are conjugate subgroups, so this gives us that all other stabilizers are likewise trivial. So,

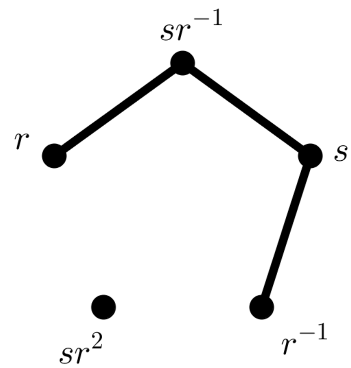

Figure 3.4: Let $S := \{r, r^{-1}, s, sr^{-1}, sr^2\}$. Above is the subgraph of $\Gamma(D_n, S)$ containing only the vertices in S and the edges among them.

to show a Cayley graph $\Gamma(G, S)$ is a GRR of its generating group, it suffices to find a generating subset R of S that is fixed pointwise by any graph automorphism fixing e. This is most straightforward when the subgraph X of $\Gamma(G, S)$ consisting of the vertices in S and the edges between them has few symmetries. An automorphism of $\Gamma(G, S)$ fixing e fixes S setwise and thus induces an automorphism of X. If X is not very symmetric, it's easier to show that the induced automorphism is always trivial.

Proposition 3.1 is used in results of both narrow and broad scope; for example, it is used to prove that D_n has a GRR for $n \geq 6$ [29], and to prove that if a group G generated by its odd-order elements has a subgroup H generated by its odd-order elements that admits a GRR, then G also has a GRR [15]. Of course, the manner in which the desired subsets are found and shown to satisfy the conditions of Proposition 3.1 are varied and frequently quite intricate. We give the proof of the dihedral result below, to give an idea of the flavor of these results.

Theorem 3.3 ([29]). *Let $n \geq 6$. Then D_n admits a graphical regular representation.*

Proof. Let $n \geq 6$ be fixed, and let $S := \{r, r^{-1}, s, sr^{-1}, sr^2\}$ be a subset of D_n. We claim $\Gamma := \Gamma(D_n, S)$ is a GRR of D_n.

Let $\phi \in A(\Gamma)_e$. Since ϕ is an automorphism and all edges involving e are of the form $\{e, s\}$ where $s \in S$, $\{\phi(e), \phi(s)\} = \{e, \phi(s)\}$ is an edge of Γ and $\phi(s)$ must be an element of S. In other words, $\phi(S) = S$. It can easily be verified that the only edges contained among vertices of S are $\{r, sr^{-1}\}$, $\{sr^{-1}, s\}$ and $\{s, r^{-1}\}$. If we restrict our attention to S and the edges contained in S, we get a path on 4 vertices and one isolated vertex (see Figure 3.4). Since ϕ fixes S setwise, $\phi|_S$ must be an automorphism of this subgraph. If $\phi|_S$ is trivial, S satisfies the conditions of Proposition 3.1, and thus Γ is a GRR.

Suppose $\phi|_S$ is non-trivial. Then ϕ must switch r and r^{-1}, switch s and sr^{-1}, and fix sr^2. Note that r^2 is adjacent to both r and s, as $r^2 = r \cdot r = s(sr^2)$. This means that $\phi(r^2)$ must be adjacent to $\phi(r) = r^{-1}$ and $\phi(s) = sr^{-1}$. In other words, there exist $t_1, t_2 \in S$ such that $\phi(r^2) = r^{-1}t_1 = sr^{-1}t_2$. We then have $t_1 = sr^{-2}t_2$. The only possible choice for t_2 is sr^{-1}, as all other choices give a t_1 which is not an element of S. But in this case $t_1 = r$, implying $\phi(r^2) = e$, a contradiction. $\square$

We end our discussion of GRRs with a conjecture of Babai, Godsil, Imrich, and Lovász, which provides an interesting contrast to the difficulty of constructing a GRR for a given group. The conjecture, stated by Babai and Godsil in [4], is that for a finite group G which is neither abelian of exponent greater than 2 nor generalized dicylic, the probability that a random Cayley graph of G is a GRR approaches 1 as $|G|$ goes to infinity. In other words, almost all Cayley graphs of a sufficiently large group G are GRRs unless G admits no GRRs at all. Babai and Godsil [4] show that this holds for nilpotent groups of odd order, but no further work has been published on the conjecture.

3.2.2 Groups with DRRs

The DRR question was settled completely in a paper by Babai [2], using the method suggested by Proposition 3.1 (which holds for digraphs as well). His result highlights the difference in difficulty in the two problems; constructing a GRR in general depended heavily on the structure of the group, but Babai provides a general construction of a DRR that works for the overwhelming majority of groups that admit a DRR. This difference is largely due to the fact that digraphs have fewer symmetries, in general, than graphs, which makes applying Proposition 3.1 easier.

We first state Babai's result, and then briefly discuss his construction.

Theorem 3.4 ([2]). *Let G be a finite group. G admits a DRR if and only if $G \notin \{\mathbb{Z}_2^2, \mathbb{Z}_2^3, \mathbb{Z}_2^4, \mathbb{Z}_3^2, Q\}$, where Q is the quaternion group.*

Note that all Cayley digraphs of $\mathbb{Z}_2^n$ are in fact graphs, so $\mathbb{Z}_2^2$, $\mathbb{Z}_2^3$, and $\mathbb{Z}_2^4$ not admitting DRRs follows from the fact that they do not admit GRRs. The proof that $\mathbb{Z}_3^2$ and Q do not admit DRRs in fact shows that they are in Class II (or, to be more precise, the analog of Watkins' Class II for digraphs), and amounts to a careful examination of possible generating sets.

Babai's general DRR construction is as follows.

Proposition 3.2. *Let G be a finite, noncyclic group with a minimal generating set $R = \{r_1, \ldots, r_d\}$ containing no elements of order two such that for some $j \in \{1, \ldots, d\}$, all of the following hold:*

i.) *if $j = 1$, $r_j^3 \neq e$;*

ii.) *if $j = d$, $r_1 r_j \neq r_j r_1$;*

iii.) *if $j = d$, $r_1^2 \neq r_j^{-2}$;*

iv.) $r_j^{-1} r_1 r_j \neq r_1^{-1}$;

v.) $r_{d-1}^{-1} r_d r_{d-1} \neq r_d^{-1}$.

Then there exists a symmetric subset $T \subseteq G$ such that $e \notin T$, $R \cap T = \emptyset$ and $\vec{\Gamma}(G, R \cup T)$ is a DRR of G.

To give the reader some idea of the proof of this proposition, the subset R here is the subset to which Babai will apply Proposition 3.1. Let $Q := \{r_i^{-1}r_{i+1} : i = 1, \ldots, d-1\} \cup \{r_j r_1\}$; T is defined as $Q \cup Q^{-1}$. T is defined solely to create a low-symmetry subgraph in the subgraph consisting of the vertices in $R \cup T$ and the edges between them. Centrally important to the proof is the fact that T is symmetric, so for all $t \in T$ and $g \in G$, both (g, gt) and $(gt, gtt^{-1}) = (gt, g)$ are in the edge set of $\vec{\Gamma}(G, R \cup T)$. All undirected edges of $\vec{\Gamma}(G, R \cup T)$ are of this form; for $r \in R$, only (g, gr) is an edge, as r^{-1} is not in R by minimality. The undirected edges among the vertices in R form an undirected path $r_1, \ldots, r_d$. If $\phi \in A(\vec{\Gamma}(G, R \cup T))$ fixes e, then $\phi(R) = R$, since R consists of exactly the vertices connected to e by directed edges coming from e. Thus, $\phi|_R$ must be an automorphism of the undirected path $r_1, \ldots, r_d$. The limited automorphisms of undirected paths and the assumptions made about R allow Babai to conclude that $\phi|R$ is trivial.

Babai's construction works for all finite groups admitting a DRR besides cyclic groups and generalized dihedral groups. A *generalized dihedral group* is a group G with an abelian subgroup A of index two such that there exists a $b \in G \setminus A$ with $b^2 = e$ and $bab = a^{-1}$ for all $a \in A$. All generating sets of a generalized dihedral group contain elements of order two, so they don't satisfy the conditions of Proposition 3.2. However, Babai shows that as long as A is cyclic (and $|A| > 2$) or satisfies the conditions of Proposition 3.2, G has a DRR. The remaining generalized dihedral groups are abelian groups of exponent 2 and order greater than 2^4 and a group of order 18; the former were shown to admit a GRR and thus a DRR by [17], and Babai proves the latter admits a DRR with little trouble. Finally, it's easy to give a DRR of a cyclic group, namely $\vec{\Gamma}(\langle a \rangle, \{a\})$, and this concludes Babai's result.

Chapter 4

Relationships between Groups and their Graphs

The previous chapters of this thesis have focused largely on one direction of the relationship between a group and its Cayley (di)graphs: roughly, what the structure of the group tells us about the structure of its Cayley (di)graphs. Chapter 2 addresses this issue at a quite basic level, exploring which group properties are "visibly" apparent in its edge-colored Cayley digraphs, while Chapter 3 investigates this relationship more subtly, discussing which groups have a Cayley (di)graph with a prescribed structure. The results in these chapters tend to use algebraic facts about a group to conclude something combinatorial about one or more of its Cayley (di)graphs. Results of this type, or in which algebraic questions about a group are translated into questions about its Cayley (di)graphs and are then solved in the Cayley (di)graph context, arise in many more topics than those presented in Chapters 2 and 3, including some in combinatorial group theory.

Less common are results involving the opposite direction of the relationship between group and Cayley (di)graph: what the combinatorial properties of a Cayley (di)graph imply about the algebraic structure of its generating group. In this chapter, we concern ourselves with a collection of results of this sort. The combinatorial property we are interested in here is proper vertex-coloring, a very well-explored and significant topic in graph theory. The question we ask about vertex-colorings of Cayley (di)graphs is motivated by how we understand Cayley (di)graphs as a class of (di)graphs. Sabidussi's Theorem (Proposition 1.3) establishes that the identifying feature of a Cayley (di)graph is its symmetry, or, speaking more specifically, that its automorphism group has a regular subgroup which is isomorphic to the Cayley (di)graph's generating group. (This regular subgroup is in fact the left regular representation $\widetilde{G}$ of the generating group G.) One particularly meaningful kind of vertex-coloring, then, would be one that somehow reflects this symmetry, perhaps by interacting in a particularly tidy way with some or all of the automorphisms in this regular subgroup. This is our avenue of investigation. We ask what it implies about the structure of a group G if one of its Cayley (di)graphs has a coloring such that certain automorphisms in $\widetilde{G}$ map color classes to themselves (our first question) or map color classes to other color classes (our second).

Before beginning, a few notes should be mentioned. We slightly modify the usual definition of a vertex coloring as a partition of the vertices into color classes. Here, we impose an ordering on the set of color classes according to the natural ordering on the colors $\{1, \ldots, n\}$. We'll write $C = [C_1, \ldots, C_n]$ to indicate that the coloring C is ordered. As there is no distinction between directed and undirected edges in the context of proper vertex colorings, we'll restrict our attention in the following sections to Cayley digraphs, with the understanding that all results also hold for Cayley graphs. In the proofs that follow, given a generating set S of a group G, we assume representatives have been chosen from each equivalence class of words in S, so each element $g \in G$ corresponds to a unique word in S.

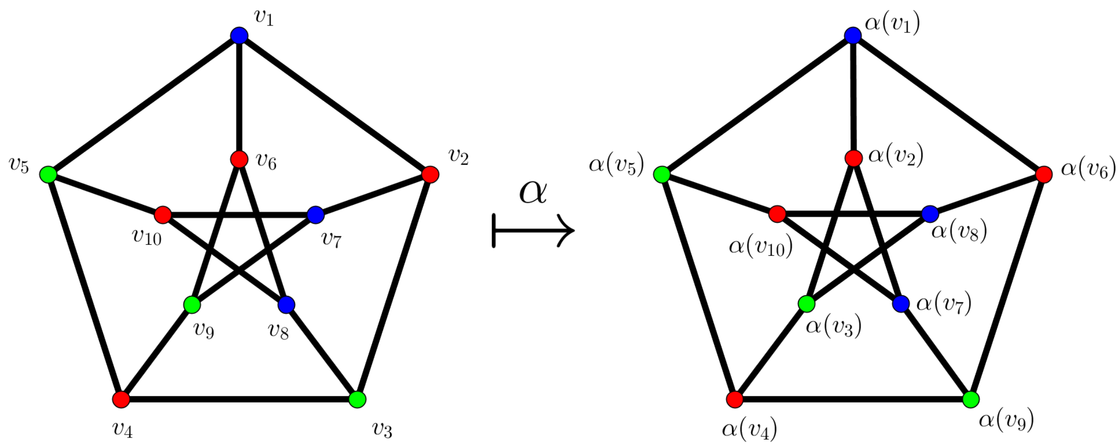

Figure 4.1: A coloring of the Petersen graph strongly compatible with the automorphism α. As you can see, the image of the colored graph under α is identical to the original.

4.1 Colorings Fixed by Automorphisms

Let $\vec{X}$ be a digraph, and $c : V(\vec{X}) \to \{1, \ldots, n\}$ some proper vertex-coloring of $\vec{X}$. In this section, our notion of an automorphism $\alpha \in A(\vec{X})$ interacting "nicely" with c is that α preserves the colors of vertices. More formally, for all $v \in V(\vec{X})$, $c(v) = c(\alpha(v))$. Alternately, if we think of the coloring c as a partition $C = [C_1, \ldots, C_n]$ of the vertices into color classes, $\alpha(C_i) = C_i$ for all $i \in \{1, \ldots, n\}$. In this case, we will say that c (or C) and α are *strongly compatible*. See Figure 4.1 for an example.

As mentioned above, for Cayley digraphs, we choose to focus on the automorphisms in $\widetilde{G}$, which are of the form $\lambda_g : x \mapsto gx$ for $g \in G$. We are particularly interested, then, in the existence of proper colorings strongly compatible with some automorphisms in $\widetilde{G}$, and what their existence or non-existence implies about the structure of G. Ideally, we would like to consider proper colorings strongly compatible with all automorphisms in $\widetilde{G}$. However, our first result shows that being strongly compatible with an automorphism is too restrictive a notion for this to be possible.

Throughout this section, G is a finite group and $S \subseteq G$ is a generating set. We are interested only in nontrivial vertex-colorings, meaning that at least one color class has size greater than one.

Proposition 4.1. *Let $e \neq g \in G$. There exists a nontrivial proper coloring $C = [C_1, \ldots, C_n]$ of $\vec{\Gamma}(G, S)$ which is strongly compatible with λ_g if and only if $h^{-1}\langle g \rangle h \cap S = \emptyset$ for all $h \in G$.*

Proof. Suppose such a coloring C exists. First, we note that color classes are unions of right cosets of $\langle g \rangle$, as $h \in C_i$ ($i \in \{1, \ldots, n\}$) implies $\lambda_g(h) = gh \in C_i$, which by the same reasoning implies $g^2h \in C_i$ and so on. This yields that $\langle g \rangle h \subseteq C_i$. As C is a proper coloring, C_i cannot contain any neighbors of h, which implies $\langle g \rangle h$ must not contain any neighbors of h. All neighbors of h are of the form hs or hs^{-1} for $s \in S$, so this statement is true if and only if $\langle g \rangle h \neq \langle g \rangle hs$ and $\langle g \rangle h \neq \langle g \rangle hs^{-1}$ for all $s \in S$. This is in turn is equivalent to $hsh^{-1} \notin \langle g \rangle$ for all $s \in S$, or slightly more concisely, $h^{-1}\langle g \rangle h \cap S = \emptyset$.

Suppose $h^{-1}\langle g \rangle h \cap S = \emptyset$ for all $h \in G$. Then color $h \in G$ by $\langle g \rangle h$. By the sequence of equivalences above, the assumption implies that for any $h_0 \in G$, no neighbors of h_0 are in $\langle g \rangle h_0$. This implies that the coloring is proper. $\square$

From a combinatorial viewpoint, it is significant that the coloring given in the proof of sufficiency is not necessarily a coloring strongly compatible with λ_g using the fewest possible colors— some color classes potentially could be combined without making the coloring improper. There is, however, no obvious way to do this in general. In other words, while the structure of some colorings strongly compatible with λ_g is easy to understand, the structure of minimal colorings having this property is not.

The following easy corollary tells us that no proper coloring can ever be strongly compatible with all automorphisms in $\widetilde{G}$.

Corollary 4.1. *Let $g \in G$. If there exists a nontrivial proper coloring of $\vec{\Gamma}(G, S)$ which is strongly compatible with λ_g, then $g \notin S$.*

In fact, for $s \in S$, λ_s is sometimes very far from being strongly compatible with colorings that are strongly compatible with other automorphisms in $\widetilde{G}$.

Proposition 4.2. *Let $g \in G$ be a nonidentity element such that $h^{-1}\langle g\rangle h \cap S = \emptyset$ for all $h \in G$. Consider the proper coloring $c : h \mapsto \langle g\rangle h$ of $\vec{\Gamma}(G, S)$. Then for all $s \in S$ and $h \in G$, $c(h) \neq c(\lambda_s(h))$. Further λ_s takes color classes to color classes if and only if s is in the normalizer of $\langle g\rangle$.*

Proof. We show the first statement by contradiction. Suppose for some $s \in S$ and $h \in G$, $c(h) = c(\lambda_s(h)) = c(sh)$. In other words, $\langle g\rangle h = \langle g\rangle sh$, which holds if and only if $s \in \langle g\rangle$, a contradiction.

The second statement follows from the fact that for $x, y \in G$, $c(x) = c(y)$ if and only if $\langle g\rangle x = \langle g\rangle y$, which occurs if and only if $xy^{-1} \in \langle g\rangle$. Suppose $c(x) = c(y)$ implies $c(sx) = c(sy)$ or, equivalently, $xy^{-1} \in \langle g\rangle$ implies $sxy^{-1}s^{-1} \in \langle g\rangle$. We rewrite $g' \in \langle g\rangle$ as $g'e^{-1}$. We have, then, that for all $g' \in \langle g\rangle$, $sg'e^{-1}s^{-1} = sg's^{-1} \in \langle g\rangle$. By definition, s is in the normalizer of $\langle g\rangle$. Suppose now that s normalizes $\langle g\rangle$. Then if $xy^{-1} \in \langle g\rangle$, $sxy^{-1}s^{-1}$ is in $\langle g\rangle$, which by the discussion above is equivalent to $c(x) = c(y)$ implying $c(sx) = c(sy)$.

□

4.2 Colorings Permuted by Automorphisms

We turn our attention to a looser notion of interaction between vertex-colorings and graph automorphisms. This notion was introduced by Chvátal and Sichler in [8], who were interested in minimal vertex colorings that reflected some symmetry of the graph being colored.

Definition 4.1. Let $\vec{X}$ be a digraph, $C = [C_1, C_2, \ldots, C_n]$ a proper vertex coloring of $\vec{X}$, and $\alpha \in A(X)$ an automorphism. C and α are *compatible* if $\alpha(C_i) \in C$

$\forall i \in \{1, \ldots, n\}$. C is *B-compatible* for $B \subseteq A(X)$ if C is compatible with β for all $\beta \in B$.

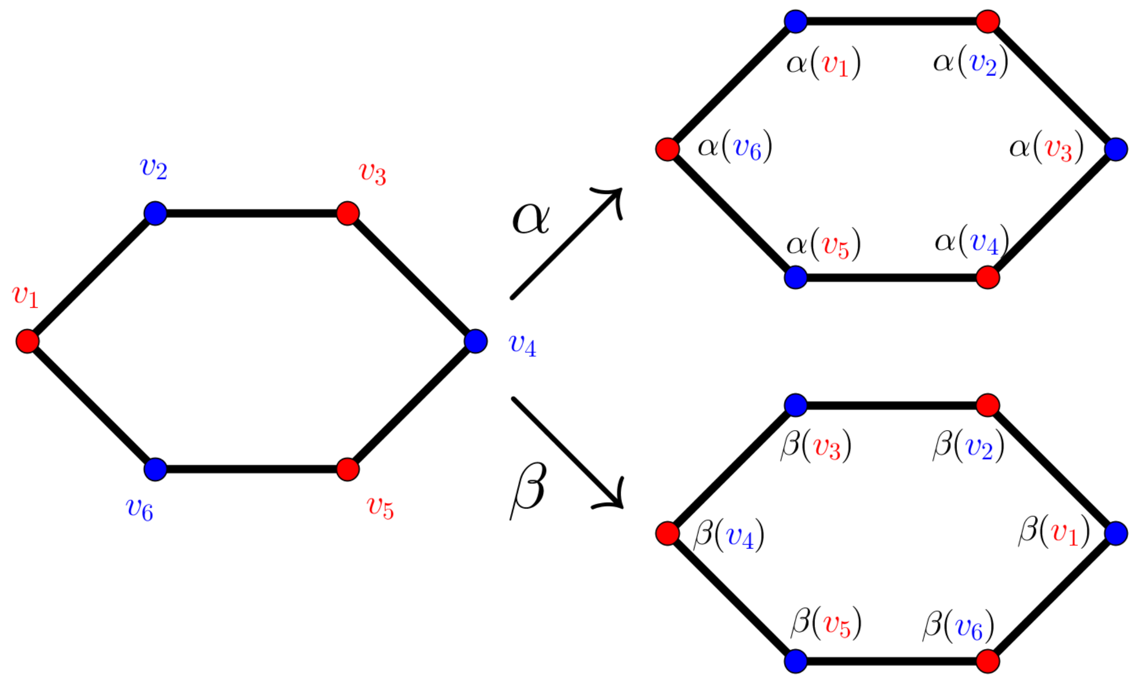

Figure 4.2: This coloring C of a cycle X on 6 vertices is compatible with the automorphisms α and β; the graph on the left shows the coloring, while the graphs on the right show that both α and β switch the color classes. As α and β generate $A(X) \cong D_6$, C is actually $A(X)$-compatible.

Less formally, a coloring and an automorphism are compatible if the automorphism takes color classes to color classes (see Figure 4.2 for an example). In fact, if a coloring C and an automorphism α are compatible, α will induce a permutation of the colors (and color classes), which follows from the fact that α is bijective on $V(\vec{X})$. Note that strongly compatible automorphisms are precisely those for which this induced permutation is the identity.

This connection between compatible automorphisms and permutations of the color classes gives rise to a number of questions. For a given coloring, what color class permutations can be realized by an graph automorphism? Or, given a permutation, can we find a proper coloring of a digraph and an compatible automorphism such that the automorphism induces the chosen permutation on the color classes? The topic of the previous section was a specific case of this question: we wanted to find a coloring of a Cayley digraph and a compatible automorphism in $\widetilde{G}$ that induced the trivial permutation on the color classes.

The focus of this section is another basic question emerging from this definition: for a given digraph automorphism (or collection of automorphisms), when does there exist a compatible coloring? More specifically, which groups have a Cayley digraph admitting a $\widetilde{G}$-compatible coloring? As in the previous section, G here is a finite group, $S \subseteq G$ a generating set and colorings are nontrivial.

Our first result gives a necessary and sufficient group theoretic condition for a Cayley digraph to have a $\widetilde{G}$-compatible coloring. The proof shows that the structure of $\widetilde{G}$-compatible colorings are related to G in a very straightforward way.

Proposition 4.3. *There exists a nontrivial proper $\widetilde{G}$-compatible coloring $C = [C_1, \ldots, C_n]$ of $\vec{\Gamma}(G,S)$ if and only if there exists $1 \neq F \leq G$ such that $F \cap S = \emptyset$.*

Proof. Suppose there exists such a coloring. In a slight abuse of notation, we denote by $C(g)$ the color class of a group element g. First, $C(e)$ is a subgroup of G, as follows. It is clearly nonempty. Let $g, h \in C(e)$. $\lambda_{gh^{-1}}(h) = g$, and as C is compatible with $\lambda_{gh^{-1}}$, this implies that $\lambda_{gh^{-1}}(C(e)) = C(e)$. Thus, $\lambda_{gh^{-1}}(e) = gh^{-1} \in C(e)$, implying $C(e) \leq G$.

Moreover, the color classes of this coloring must be left cosets of $C(e)$. Let $g \in G$. Then $\lambda_g(e) = g$, which implies $\lambda_g(C(e)) = C(g)$. On the other hand, $\lambda_g(C(e)) = gC(e)$, so we have $gC(e) = C(g)$.

As C is a proper coloring, for all $g \in G$, no neighbors of g can be in $C(g) = gC(e)$. The neighbors of g are of the form gs and gs^{-1}, where $s \in S$, so g having no neighbors of the same color is equivalent to $gC(e) \neq (gs)C(e)$ and $gC(e) \neq (gs^{-1})C(e)$ for all $s \in S$. This implies that $g^{-1}gs = s \notin C(e)$. Additionally, as C is nontrivial, $C(e) \neq 1$, so $C(e)$ is a subgroup of the kind we sought.

Now suppose there exists such a subgroup F. Then color $g \in G$ by gF. By $F \cap S = \emptyset$, $gF \neq (gs)F$ for all $s \in S$. We also have that $gF \neq (gs^{-1})F$ for all $s \in S$, as F is closed under inverses. This means that no neighbors of g are in its color class, so the coloring is proper. It is also compatible with the desired automorphisms, as $\lambda_g(hF) = ghF = (gh)F$, and left cosets are either equal or disjoint.

$\square$

The proof of this result gives us plentiful information about the structure of $\widetilde{G}$-compatible colorings, when they exist. Namely, the color classes are cosets of a subgroup avoiding the generating set. This means the color classes all contain the same number of vertices, so the number of colors used is equal to the index of that subgroup. The proof of sufficiency also tells us that for each subgroup avoiding the generating set S (and some fixed ordering of its cosets), there is one $\widetilde{G}$-compatible coloring of $\vec{\Gamma}(G, S)$.

Proposition 4.3 turns the question of a Cayley digraph having a $\widetilde{G}$-compatible coloring into a purely group theoretic question. Which groups have generating sets avoiding at least one of its subgroups? There's an easy strategy to find a generating set which doesn't avoid any subgroups: simply take a minimal generating set and add one element from each subgroup. A wise choice of these additional elements will generally yield a generating set which is not the whole group. This motivates us to restrict our attention to minimal generating sets in the following discussion.

Subgroups which avoid every minimal generating set of a group are a well-investigated object in group theory. The *Frattini subgroup* of a group G, denoted $\Phi(G)$, is defined as the intersection of the maximal subgroups of G (or G if this intersection is empty). $\Phi(G)$ is precisely the set of elements of G which cannot appear in a minimal generating set [23, pg. 135]. If $1 \neq \Phi(G) \lneq G$, then G has a nontrivial subgroup avoiding every minimal generating set. Thus the coloring whose color classes are cosets of $\Phi(G)$ is a proper $\widetilde{G}$-compatible coloring of every minimal Cayley digraph of G.

This is quite a bit stronger than the condition we are interested in, which is the existence of a $\widetilde{G}$-compatible coloring in a particular minimal Cayley digraph $\vec{\Gamma}(G,S)$ or, equivalently, the existence of a nontrivial subgroup of G which avoids a particular minimal generating set S. Of course, groups with nontrivial proper Frattini subgroups fulfill these conditions, but a number of groups without nontrivial proper Frattini subgroups do as well. In fact, groups whose minimal generating sets intersect every nontrivial subgroup have a very specific structure, which is investigated below. We first point out a statement which is easily equivalent to the existence of a nontrivial subgroup avoiding a generating set, which will simplify later arguments.

Lemma 4.1. *There exists $1 \neq F \leq G$ such that $F \cap S = \emptyset$ if and only if there exists $e \neq g \in G$ such that $\langle g \rangle \cap S = \emptyset$.*

Proof. This is clear from the fact that for all $f \in F$, $\langle f \rangle \subseteq F$ and that $\langle g \rangle$ is a nontrivial subgroup if $e \neq g$. □

Proposition 4.4. *Let G be a finite group and $S \subseteq G$ a minimal generating set. Suppose that for all $1 \neq F \leq G$, $F \cap S$ is nonempty. Then the following conditions hold.*

i.) For all $s \in S$, $o(s)$ is prime;

ii.) Let $g \in G$. If $o(g)$ is prime, then $g \in \bigcup_{s \in S} \langle s \rangle$;

iii.) $\langle s \rangle \trianglelefteq G$ for all $s \in S$.

Proof. i.) We proceed by contrapositive. Let $s \in S$ and suppose that an integer $m \notin \{1, o(s)\}$ divides $o(s)$. Then $s \notin \langle s^m \rangle$, as $|\langle s^m \rangle| < |\langle s \rangle|$. By the minimality of S, this implies that $\langle s^m \rangle \cap S = \emptyset$, and by the choice of m, $\langle s^m \rangle$ is a nontrivial subgroup of G.

ii.) Suppose for the sake of contradiction that $o(g)$ is prime and g is not a power of any element of S. By Lemma 4.1, $\langle g \rangle \cap S$ is nonempty. In other words, for some $s \in S$ and $m \in \mathbb{Z}$, $g^m = s$. This implies $\langle g^m \rangle = \langle s \rangle$. Note that $s \neq e$, so $\langle g^m \rangle \neq 1$. Since $o(g)$ is prime, we also have that $\langle g^m \rangle = \langle g \rangle$. Thus, $g \in \langle s \rangle$, a contradiction.

iii.) Let $s \in S$. To show $\langle s \rangle$ is normal, it suffices to consider conjugates of s, as the conjugate of a power is simply a power of conjugates. We think of the conjugator as a word in S and induct on the length of this word. For the base case, let $s_1 \in S$. $o(s_1^{-1} s s_1) = o(s)$, which is prime by i). By ii), $s_1^{-1} s s_1 = \tilde{s}^m$ for some $\tilde{s} \in S$, $m \in \mathbb{Z}$. If $\tilde{s} \neq s$, then $s = s_1 \tilde{s}^m s_1^{-1}$, where the right hand side is an expression in S not involving s, contradicting the minimality of S. So $\tilde{s} = s$, and $s_1^{-1} s s_1 \in \langle s \rangle$.

Now, suppose conjugates of s by elements of word length $n-1$ are elements of $\langle s \rangle$. Consider $g \in G$ such that $g = s_{\nu_1} s_{\nu_2} \cdots s_{\nu_n}$ (note that the s_{ν_i}'s are not necessarily distinct). Then

$$\begin{aligned} g^{-1} s g &= s_{\nu_n}^{-1} \cdots s_{\nu_1}^{-1} s s_{\nu_1} \cdots s_{\nu_n} \\ &= s_{\nu_n}^{-1} (s_{\nu_{n-1}}^{-1} \cdots s_{\nu_1}^{-1} s s_{\nu_1} \cdots s_{\nu_{n-1}}) s_{\nu_n} \\ &= s_{\nu_n}^{-1} s^m s_{\nu_n} \end{aligned}$$

for some $m \in \mathbb{Z}$ by the inductive hypothesis. $s_{\nu_n}^{-1} s^m s_{\nu_n} = (s_{\nu_n}^{-1} s s_{\nu_n})^m$, and $s_{\nu_n}^{-1} s s_{\nu_n} \in \langle s \rangle$ by the base case, so we conclude that $g^{-1} s g = (s_{\nu_n}^{-1} s s_{\nu_n})^m \in \langle s \rangle$.

□

The conclusions of Proposition 4.4 are quite strong. In fact, they are enough to completely determine the structure of the groups that have minimal Cayley digraphs without $\widetilde{G}$-compatible colorings.

Proposition 4.5. *Suppose G is a finite group and S a minimal generating set of G such that $\vec{\Gamma}(G,S)$ does not have a nontrivial proper $\widetilde{G}$-compatible coloring. Then G is the direct product of cyclic groups of prime order, and S consists of exactly one element from each factor.*

Proof. By Proposition 4.3, the assumption implies that there are no nontrivial subgroups $F \leq G$ such that $F \cap S = \emptyset$, so the conclusions of Proposition 4.4 hold for G.

We induct on the size of S. If S consists of a single element, G is cyclic and Proposition 4.4 implies the order of the single generator is prime.

Now, suppose the statement holds for generating sets of size $t-1$, and suppose the assumptions hold for a group G and a minimal generating set $S = \{s_1, \ldots, s_t\}$. Let $S' := S \setminus \{s_1\}$, and let $G' := \langle S' \rangle$. We claim that $G \cong G' \times \langle s_1 \rangle$.

Indeed, $G = G'\langle s_1 \rangle$. To show this, we think of a group element g as a word in S and induct on the number of times s_1 appears in this word Suppose $g = s_{\nu_1} \ldots s_{\nu_m} \in G$ and $s_{\nu_j} = s_1$ for precisely one $j \in \{1, \ldots, m\}$. If $j = m$, the conclusion follows, so suppose $j < m$. Let $h = s_{\nu_{j+1}} \ldots s_{\nu_m}$. Since $\langle s_1 \rangle$ is normal by Proposition 4.4, $h^{-1}s_1h = s_1^n$ for some $l \in \mathbb{Z}$. This can be restated as $s_1h = hs_1^n$. So $g = s_{\nu_1} \ldots s_{\nu_{j-1}} s_{\nu_{j+1}} \ldots s_{\nu_m} s_1^n \in G'\langle s_1 \rangle$, as $s_{\nu_i} \neq s_1$ for $i \neq j$. Now, suppose that if s_1 appears in g's word representation k times ($k > 1$), then $g \in G'\langle s_1 \rangle$. Let $g = s_{\nu_1} \ldots s_{\nu_m}$ be an element of G where s_1 appears $k+1$ times in its word representation. If $s_{\alpha_m} = s_1$, then gs_1^{-1} is an element of $G'\langle s_1 \rangle$ by the inductive hypothesis, implying that g is also. Otherwise, choose $j \in \{1, \ldots, m\}$ such that $s_{\nu_j} = s_1$ and $s_{\nu_i} \neq s_1$ for $i > j$. By the base case, $s_{\nu_j} \ldots s_{\nu_m} = s_{\nu_{j+1}} \ldots s_{\nu_m} s_1^n$ for some $n \in \mathbb{Z}$, so $g = s_{\nu_1} \ldots s_{\nu_{j-1}} s_{\nu_{j+1}} \ldots s_{\nu_m} s_1^n$. s_1 appears k times in $gs_1^{-n} = s_{\nu_1} \ldots s_{\nu_{j-1}} s_{\nu_{j+1}} \ldots s_{\nu_m}$, so applying the inductive hypothesis to gs_1^{-n} yields that $g \in G'\langle s_1 \rangle$.

Further, $\langle s_1 \rangle \cap G' = 1$. Suppose, for the sake of contradiction, that $s_1^n \in G'$ for some nontrivial power n of s_1. Then $\langle s_1^n \rangle \subseteq G'$. Since s_1 is of prime order by

Proposition 4.4, $\langle s_1 \rangle = \langle s_1^n \rangle$, which implies $s_1 \in G'$, a contradiction.

Finally, $\langle s_1 \rangle, G' \trianglelefteq G$. The former follows from Proposition 4.4. To see the latter, let $h = s_{\nu_1} \dots s_{\nu_m} s_1^n \in G$ and $g' = s_{\mu_1} \dots s_{\mu_l} \in G'$, where $s_{\nu_i}, s_{\mu_j} \in S'$ for all i, j. We insert $s_1^n s_1^{-1}$ between consecutive elements of S' in $h^{-1} g' h$ to rewrite $h^{-1} g' h$ as a product of conjugates of powers of $s_k \in S'$ by s_1^n . Each of these conjugates is contained in $\langle s_k \rangle$, and so their product is contained in G'. Thus, we have that $G \cong G' \times \langle s_1 \rangle$.

Note that G' and S' also fulfill the conditions of the proposition. If S' were not a minimal generating set of G', S would not be a minimal generating set of G. If G' contained a subgroup H' disjoint from S', then G would contain a subgroup (isomorphic to $H' \times \{e\}$) disjoint from S. As S' contains $t-1$ elements, by the inductive hypothesis, $G' \cong \langle s_2 \rangle \times \cdots \times \langle s_t \rangle$. The proposition follows.

$\square$

This proposition gives us a necessary condition for a group G and a minimal generating set S such that $\vec{\Gamma}(G, S)$ does not have a $\widetilde{G}$-compatible coloring. It's natural to ask if this condition is also sufficient, and not difficult to see that it is not. See Figure 4.3 for an example of a Cayley digraph fulfilling the necessary condition, shown with a vertex coloring compatible with $\widetilde{G}$.

Proposition 4.6. *Let $G \cong \mathbb{Z}_{p_1} \times \mathbb{Z}_{p_2} \times \cdots \times \mathbb{Z}_{p_n}$, where p_i is prime and for some $i \neq j \in \{1, \dots, n\}$, $p_i = p_j$. Then for any minimal generating set S of G, $\vec{\Gamma}(G, S)$ has a nontrivial proper $\widetilde{G}$-compatible coloring.*

Proof. For simplicity, we assume $G = \mathbb{Z}_{p_1} \times \mathbb{Z}_{p_2} \times \cdots \times \mathbb{Z}_{p_n}$ and $i = 1$ and $j = 2$. Let $p := p_1$.

Suppose, for the sake of contradiction, that for some minimal generating set S of G, $\vec{\Gamma}(G, S)$ has no nontrivial proper $\widetilde{G}$-compatible coloring. By Proposition 4.5, G is the direct product of the cyclic subgroups generated by elements of S. As the

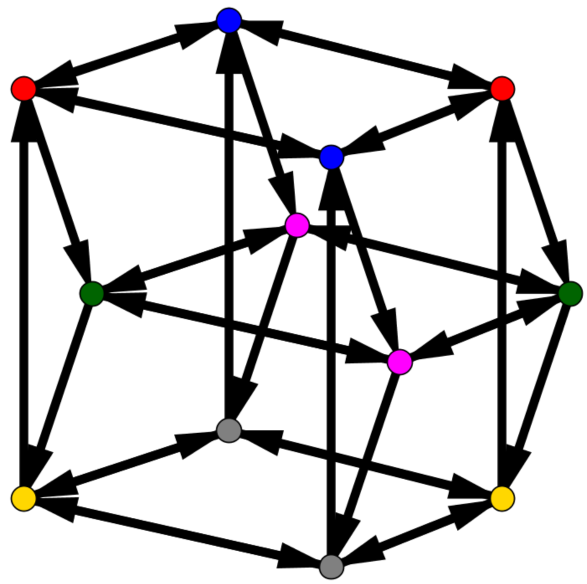

Figure 4.3: Let $G := \mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_3$, and $S := \{(\overline{1}, \overline{0}, \overline{0}), (\overline{0}, \overline{1}, \overline{0}), (\overline{0}, \overline{0}, \overline{1})\}$. $\vec{\Gamma}(G, S)$ illustrates that the necessary condition given in Proposition 4.5 is not sufficient, as it fulfills the condition but has a $\widetilde{G}$-compatible coloring, which is shown here. Note that to verify a coloring is $\widetilde{G}$-compatible, it suffices to check it is compatible with λ_s for $s \in S$, which is clearly the case here.

decomposition of a finite abelian group into the direct product of cyclic groups is unique up to ordering, S consists of one element from each $\mathbb{Z}_{p_k}$ or rather, from each subgroup of G isomorphic to $\mathbb{Z}_{p_k}$. In other words, S consists of n-tuples where all but one coordinate is $\overline{0}$. Let $g := (\overline{1}, \overline{1}, \overline{0}, \ldots, \overline{0})$. $o(g) = p$, but g is not a power of any element of S. This contradicts part ii. of Proposition 4.4. $\square$

However, direct products of cyclic groups of distinct prime orders do have Cayley digraphs without $\widetilde{G}$-compatible colorings.

Proposition 4.7. *Let $G \cong \mathbb{Z}_{p_1} \times \mathbb{Z}_{p_2} \times \cdots \times \mathbb{Z}_{p_n}$, where $p_1 < p_2 < \cdots < p_n$ are prime, and let S consist of 1 element of order p_i for all $i \in \{1, \ldots, n\}$. Then $\vec{\Gamma}(G, S)$ does not have a nontrivial proper $\widetilde{G}$-compatible coloring.*

Proof. By Proposition 4.1, it suffices to show that for all $e \neq g \in G$, $\langle g \rangle$ intersects S.

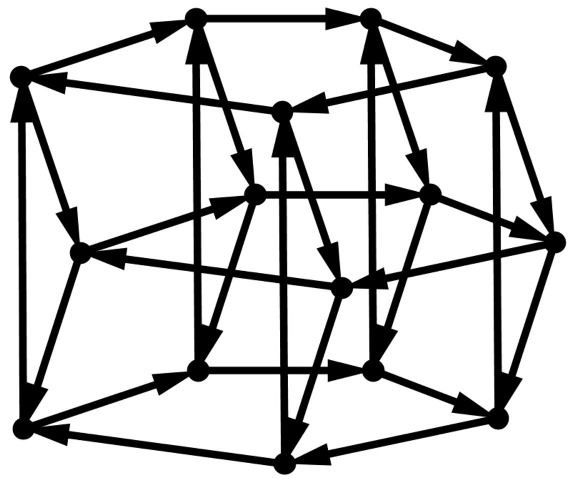

Figure 4.4: $\vec{\Gamma}(\mathbb{Z}_5 \times \mathbb{Z}_3, \{(\overline{1}, \overline{0}), (\overline{0}, \overline{1})\})$. An example of a Cayley digraph without a nontrivial $\widetilde{G}$-compatible coloring.

For simplicity, we assume $G = \mathbb{Z}_{p_1} \times \mathbb{Z}_{p_2} \times \cdots \times \mathbb{Z}_{p_n}$. By Proposition 4.5 and the same argument as in Proposition 4.3, S consists of elements with precisely one nonidentity coordinate, containing one element from $\{\overline{0}\} \times \cdots \times \mathbb{Z}_{p_k} \times \cdots \times \{\overline{0}\}$ for each $k \in \{1, \ldots, n\}$. We note that if $h \in G$ and $o(h) = p_k$ is prime, $\langle h \rangle \cap S$ is nonempty: h is an element of $\{\overline{0}\} \times \cdots \times \mathbb{Z}_{p_k} \times \cdots \times \{\overline{0}\}$, and, since this subgroup is cyclic and of prime order, $\langle h \rangle = \{\overline{0}\} \times \cdots \times \mathbb{Z}_{p_k} \times \cdots \times \{\overline{0}\}$.

Let $e \neq g \in G$. We can find $m \in \mathbb{Z}$ such that $o(g^m)$ is prime: if $o(g)$ is prime, $m := 1$. Otherwise, if some prime p divides $o(g)$, $m := \frac{o(g)}{p}$. By the observation above, $\langle g^m \rangle \cap S$ is nonempty. As $\langle g^m \rangle \subseteq \langle g \rangle$, this implies $\langle g \rangle \cap S$ is nonempty. $\square$

Propositions 4.5, 4.6, and 4.7 allow us to state another, more constructive, necessary and sufficient condition for $\vec{\Gamma}(G, S)$ to have a $\widetilde{G}$-compatible coloring.

Proposition 4.8. *Let G be a finite group and $S \subseteq$ a minimal generating set. Then $\vec{\Gamma}(G, S)$ does not have a nontrivial proper $\widetilde{G}$-compatible coloring if and only if $G \cong \mathbb{Z}_{p_1} \times \mathbb{Z}_{p_2} \times \cdots \times \mathbb{Z}_{p_n}$, where $p_1 < p_2 < \cdots < p_n$ are prime, and S consists of 1 element of order p_i for all $i \in \{1, \ldots, n\}$.*

At this point, we recall that, combinatorially speaking, we are especially interested in vertex-colorings using the fewest possible colors, called *minimal* colorings. We've ascertained that most minimal Cayley digraphs have $\widetilde{G}$-compatible colorings. However, the color classes of colorings guaranteed by various propositions are generally cosets of $\langle g \rangle$ for some group element g. The number of colors used is the index of $\langle g \rangle$, which tends to be large compared to the index of other subgroups. It seems plausible that many groups have subgroups of smaller index avoiding a given generating set S, which would yield a $\widetilde{G}$-compatible coloring of $\vec{\Gamma}(G, S)$ using fewer colors.

A new set of questions is suggested by this line of reasoning. How few colors can a $\widetilde{G}$-compatible coloring of $\vec{\Gamma}(G, S)$ use, and how does this number relate to the properties of G and S? The vertex-colorings studied in graph theory are without symmetry requirements, so another avenue of inquiry is how the minimum number of colors needed to $\widetilde{G}$-compatibly color a Cayley digraph compares with the minimum number of colors needed to just properly color the Cayley digraph (denoted $\chi(\vec{\Gamma}(G, S))$). Which Cayley digraphs $\vec{\Gamma}(G, S)$ have a $\widetilde{G}$-compatible coloring using $\chi(\vec{\Gamma}(G, S))$ colors?

The final question is equivalent to asking which groups G have generating sets S such that a subgroup of index $\chi(\vec{\Gamma}(G, S))$ avoids S. A few obvious groups and generating sets fulfill this condition. S_n with generating set $T := \{(1, 2), (1, 3), \ldots, (1, n)\}$ of transpositions is one example, as A_n is a subgroup of index 2 avoiding T and $\chi(\vec{\Gamma}(S_n, T)) = 2$. See Figure 4.5 for an example of a minimal and $\widetilde{G}$-compatible coloring of $\vec{\Gamma}(S_n, T)$. D_{2n} with the standard generating set also fulfills the condition, as $\langle r^2, sr \rangle$ is a subgroup of index 2 disjoint from $\{r, s\}$. The minimal and $\widetilde{G}$-compatible coloring this gives of $\vec{\Gamma}(D_{2n}, \{r, s\})$ is shown in Figure 4.6. Besides these examples, however, the question appears to be open, as are those in the previous paragraph.

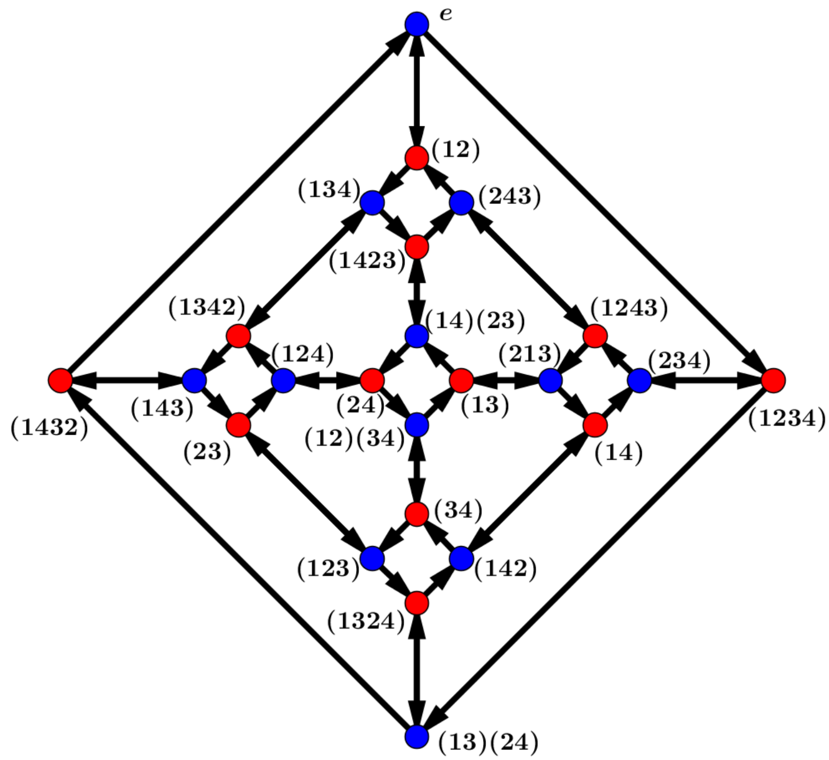

Figure 4.5: A minimal and $\widetilde{S}_4$-compatible coloring of $\vec{\Gamma}(S_4, T)$, where $T := \{(12), (1234)\}$. Color classes are cosets of A_4.

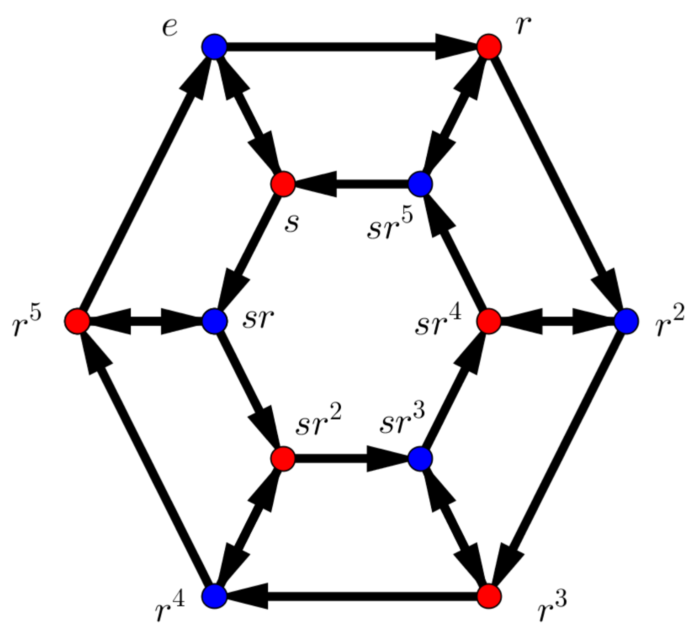

Figure 4.6: A minimal and $\widetilde{D_6}$-compatible coloring of $\vec{\Gamma}(D_6, \{r, s\})$. Color classes are cosets of $\langle r^2, sr \rangle$.

Bibliography

[1] L. Babai. On the minimum order of graphs with given group. *Canad. Math. Bull.*, 17:467–470, 1974.

[2] L. Babai. Finite digraphs with given regular automorphism groups. *Periodica Math. Hung.*, 11:257–270, 1980.

[3] L. Babai. Automorphism Groups, Isomorphisms, Reconstruction. In R. L. Graham, M. Grotschel, and L. Lovász, editors, *Handbook of Combinatorics*, volume 2, pages 1447–1541. North Holland, 1995.

[4] L. Babai and C. Godsil. On the automorphism groups of almost all Cayley graphs. *Europ. J. Combinatorics*, 3:9–15, 1982.

[5] N. Carter. *Visual Group Theory.* MAA, Washington, D. C., 2009.

[6] C. Y. Chao. On a theorem of Sabidussi. *Proc. Amer. Math. Soc.*, 15:291–292, 1964.

[7] G. Chartrand, L. Lesniak, and P. Zhang. *Graphs & Digraphs.* A Chapman & Hall book. Taylor & Francis, 5th edition, 2010.

[8] V. Chvátal and J. Sichler. Chromatic automorphisms of graphs. *J. Combinatorial Theory (B)*, 14:209–215, 1973.

[9] A. Clark. *Elements of Abstract Algebra.* Dover Books on Mathematics. Dover, New York, 1984.

[10] R. Diestel. *Graph Theory.* Electronic library of mathematics. Springer, 4th edition, 2010.

[11] D. S. Dummit and R. M. Foote. *Abstract Algebra.* John Wiley & Sons, Hoboken, NJ, 3rd edition, 2004.

[12] R. Frucht. Herstellung von Graphen mit vorgegebener abstrakter Gruppe. *Compositio Math.*, 6:239–250, 1938.

[13] R. Frucht. Graphs of degree three with a given abstract group. *Canad. J. Math*, 1:365–378, 1949.

[14] J. A. Gallian. *Contemporary Abstract Algebra.* Brooks/Cole, Boston, MA, 2013.

[15] C. D. Godsil. GRR's for non-solvable groups. In L. Lovász, editor, *Algebraic Methods in Graph Theory (Proc. Conf. Szeged, 1978)*, pages 221–239. North-Holland, 1981.

[16] D. Hetzel. *Über reguläre graphische Darstellung von auflösbaren Grupen.* PhD thesis, Technische Universität Berlin, 1976.

[17] W. Imrich. Graphs with transitive abelian automorphism group. In P. Erdős, A. Renyi, and V. T. Sós, editors, *Combinatorial Theory and Its Applications (Coll. Math. Soc. János Bolyai 4, Balatonfüred, Hungary, 1969)*, pages 651–656. North-Holland, 1970.

[18] A. Karrass, W. Magnus, and D. Solitar. *Combinatorial Group Theory.* Dover, New York, 2nd edition, 1976.

[19] D. König. *Theorie der endlichen und unendlichen Graphen.* Akad. Verlagsges. Geest u. Portig, Liepzig, 1936.

[20] L. A. Nowitz. On the non-existence of graphs with transitive generalized dicyclic groups. *J. Comb. Theory*, 4:49–51, 1968.

[21] L. A. Nowitz and M. E. Watkins. Graphical regular representations of non-abelian groups, I. *Can. J. Math*, 24:993–1008, 1972.

[22] L. A. Nowitz and M. E. Watkins. Graphical regular representations of non-abelian groups, II. *Can. J. Math*, 24:1009–1018, 1972.

[23] D. J. S. Robinson. *A Course in the Theory of Groups*, volume 80 of *Graduate Texts in Mathematics*. Springer, New York, 2nd edition, 1996.

[24] G. Sabidussi. On a class of fixed-point-free graphs. *Proc. Amer. Math. Soc.*, 9:800–804, 1958.

[25] G. Sabidussi. On the minimum order of graphs with a given automorphism group. *Montash. Math*, 63:124–127, 1959.

[26] T. Tao. *Cayley graphs and the algebra of groups*, 2012 (accessed August 11, 2015). `https://terrytao.wordpress.com/2012/05/11/cayley-graphs-and-the-algebra-of-groups/`.

[27] M. E. Watkins. On graphical regular representations of $C_n \times Q$. In *Graph theory and its applications*, pages 305–311. Springer, 1972.

[28] M. E. Watkins. Graphical regular representation of alternating, symmetric and miscellaneous small groups. *Aequat. Math.*, 11:40–50, 1974.

[29] M.E. Watkins. On the action of non-abelian groups on graphs. *J. Comb. Theory B*, 1:95–104, 1971.