

Cambridge University Press

978-0-521-89545-3 - Groups, Graphs and Trees: An Introduction to the Geometry of Infinite Groups

John Meier

Excerpt

[More information](#)

1

Cayley's Theorems

As for everything else, so for a mathematical theory: beauty can be perceived but not explained.

–Arthur Cayley

An introduction to group theory often begins with a number of examples of finite groups (symmetric, alternating, dihedral, ...) and constructions for combining groups into larger groups (direct products, for example). Then one encounters Cayley's Theorem, claiming that every finite group can be viewed as a subgroup of a symmetric group. This chapter begins by recalling Cayley's Theorem, then establishes notation, terminology, and background material, and concludes with the construction and elementary exploration of Cayley graphs. This is the foundation we use throughout the rest of the text where we present a series of variations on Cayley's original insight that are particularly appropriate for the study of infinite groups.

Relative to the rest of the text, this chapter is gentle, and should contain material that is somewhat familiar to the reader. A reader who has not previously studied groups and encountered graphs will find the treatment presented here "brisk."

1.1 Cayley's Basic Theorem

You probably already have good intuition for what it means for a group to act on a set or geometric object. For example:

- The cyclic group of order n – denoted \mathbb{Z}_n – acts by rotations on a regular n -sided polygon.

Cambridge University Press

978-0-521-89545-3 - Groups, Graphs and Trees: An Introduction to the Geometry of Infinite Groups

John Meier

Excerpt

[More information](#)

2

Cayley's Theorems

- The dihedral group of order $2n$ – denoted D_n – also acts on the regular n -sided polygon, where the elements either rotate or reflect the polygon.
- We use SYM_n to denote the symmetric group of all permutations of $[n] = \{1, 2, \dots, n\}$. (More common notations are S_n and Σ_n .) By its definition, SYM_n acts on this set of numbers, as does its index 2 subgroup, the alternating group A_n , consisting of the even permutations.
- Matrix groups, such $\text{GL}_n(\mathbb{R})$ (the group of invertible n -by- n matrices with real number entries), act on vector spaces.

Because the general theme of this book is to study groups via actions, we need a bit of notation and a formal definition.

Convention 1.1. If X is a mathematical object (such as a regular polygon or a set of numbers), then we use $\text{SYM}(X)$ to denote all bijections from X to X that preserve the indicated mathematical structure. For example, if X is a set, then $\text{SYM}(X)$ is simply the group of permutations of the elements of X . In fact, if $n = |X|$ then $\text{SYM}(X) \approx \text{SYM}_n$. Moreover, if X and X' have the same cardinality, then $\text{SYM}(X) \approx \text{SYM}(X')$. If X is a regular polygon, then angles and lengths are important, and $\text{SYM}(X)$ will be composed of rotations and reflections (and it will in fact be a dihedral group). Similarly, if X is a vector space, then $\text{SYM}(X)$ will consist of bijective linear transformations.

What we are referring to as “ $\text{SYM}(X)$ ” does have a number of different names in different contexts within mathematics. For example, if G is a group, then the collection of its symmetries is referred to as $\text{AUT}(G)$, the group of automorphisms. If we are working with the Euclidean plane, \mathbb{R}^2 , and are considering functions that preserve the distance between points, then we are looking at $\text{ISOM}(\mathbb{R}^2)$, the group of isometries of the plane.

It is quite useful to have individual names for these groups, as their names highlight what mathematical structures are being preserved. Our convention of lumping these various groups all together under the name “ SYM ” is vague, but we believe that in context it will be clear what is intended, and we like the fact that this uniform terminology emphasizes that these various situations where groups arise are not all that different.¹ One egregious example, which highlights the need to be care-

¹ In his book, *Symmetry*, Hermann Weyl wrote: “[W]hat has indeed become a guiding principle in modern mathematics is this lesson: *Whenever you have to deal with a structure-endowed entity Σ try to determine its group of automorphisms, the group of those element-wise transformations which leave all structural relations*

Cambridge University Press

978-0-521-89545-3 - Groups, Graphs and Trees: An Introduction to the Geometry of Infinite Groups

John Meier

Excerpt

[More information](#)

1.1 Cayley's Basic Theorem

3

ful in using our convention, comes from the integers. If the integers are thought of as simply a set, containing infinitely many elements, then $\text{SYM}(\mathbb{Z})$ is an infinite permutation group, which contains SYM_n for any n . On the other hand, if \mathbb{Z} denotes the group of integers under addition, then $\text{SYM}(\mathbb{Z}) \approx \mathbb{Z}_2$. (The only non-trivial automorphism of the group of integers sends n to $-n$ for all $n \in \mathbb{Z}$.)

Definition 1.2. An *action* of a group G on a mathematical object X is a group homomorphism from G to $\text{SYM}(X)$. Equivalently, it is a map from $G \times X \rightarrow X$ such that

1. $e \cdot x = x$, for all $x \in X$; and
2. $(gh) \cdot x = g \cdot (h \cdot x)$, for all $g, h \in G$ and $x \in X$.

We denote “ G acts on X ” by $G \curvearrowright X$.

If one has a group action $G \curvearrowright X$, then the associated homomorphism is a *representation* of G . The representation is *faithful* if the map is injective. In other words, it is faithful if, given any non-identity element $g \in G$, there is some $x \in X$ such that $g \cdot x \neq x$.

Example 1.3. The dihedral group D_n is the symmetry group of a regular n -gon. As such, it also permutes the vertices of the n -gon, hence there is a representation $D_n \rightarrow \text{SYM}_n$. As every non-identity element of D_n moves at least $(n - 2)$ vertices, this representation is faithful.

Remark 1.4 (left vs. right). In terms of avoiding confusion, this is perhaps the most important remark in this book. Because not all groups are abelian, it is very important to keep left and right straight. All of our actions will be *left* actions (as described above). We have chosen to work with left actions since it matches function notation and because left actions are standard in geometric group theory and topology.

Groups arise in a number of different contexts, most commonly as symmetries of any one of a number of possible mathematical objects X . In these situations, one can often understand the group directly from our understanding of X . The dihedral and symmetric groups are two examples of this. However, groups are abstract objects, being merely a set with a binary operation that satisfies a certain minimal list of requirements. Cayley's Theorem shows that the abstract notion of a group and the notion of a group of permutations are one and the same.

undisturbed. You can expect to gain a deep insight into the constitution of Σ in this way.” Our use of $\text{SYM}(\Sigma)$ instead of $\text{AUT}(\Sigma)$ is a small notational deviation from Weyl's recommendation.

Cambridge University Press

978-0-521-89545-3 - Groups, Graphs and Trees: An Introduction to the Geometry of Infinite Groups

John Meier

Excerpt

[More information](#)

Theorem 1.5 (Cayley's Basic Theorem). *Every group can be faithfully represented as a group of permutations.*

Proof. The objects that G permutes are the elements of G . In this proof we use " SYM_G " to denote $\text{SYM}(G)$, to emphasize that " G " denotes the underlying set of elements, not the group. The permutation associated to $g \in G$ is defined by left multiplication by g . That is, $g \mapsto \pi_g \in \text{SYM}_G$ where $\pi_g(h) = g \cdot h$ for all $h \in G$. This is a permutation of the elements of G , since if $g \cdot h = g \cdot h'$, then by left cancellation, $h = h'$. Denote the map taking the element g to the permutation π_g by $\pi : G \rightarrow \text{SYM}_G$.

To check that π is a group homomorphism we need to verify that $\pi(gh) = \pi(g) \cdot \pi(h)$. In other words, we need to show that $\pi_{gh} = \pi_g \cdot \pi_h$. We do this by evaluating what each side does to an arbitrary element of G . We denote the arbitrary element by " x ", thinking of it as a variable. The permutation π_{gh} takes $x \mapsto (gh) \cdot x$, and successively applying π_h then π_g sends $x \mapsto h \cdot x \mapsto g \cdot (h \cdot x)$. Thus checking that ϕ is a homomorphism amounts to verifying the associative law: $(gh) \cdot x = g \cdot (h \cdot x)$. As this is part of the definition of a group, the equation holds.

In order to see that the map is faithful it suffices to show that no non-identity element is mapped to the trivial permutation. One can do this by simply noting that if $g \in G \setminus \{e\}$, then $g \cdot e = g$, hence $\pi_g(e) = g$, and so π_g is not the identity (or trivial) permutation. \square

The proof of Cayley's Basic Theorem constructs a representation of G as a group of permutations of itself. Before moving on we should examine what these permutations look like in some concrete situations. We first consider SYM_3 , the group of all permutations of three objects.

Notation 1.6 (cycle notation). In describing elements of SYM_n we use cycle notation, and multiply (that is, compose permutations) right to left. This matches with our intuition from functions where $f \circ g(x)$ means that you first apply g then apply f , and it is consistent with our use of left actions. Here is a concrete example: $(12)(35) \in \text{SYM}_5$ is the element that transposes 1 and 2, as well as 3 and 5; the element (234) sends 2 to 3, 3 to 4 and 4 to 2; the product $(12)(35) \cdot (234) = (12534)$. (The product is not (13542) , which is the result of multiplying left to right.)

Example 1.7. The group SYM_3 has six elements, shown as disjoint vertices in Figure 1.1. The permutations described by Cayley's Basic Theorem – for the elements (12) and (123) – are also shown.

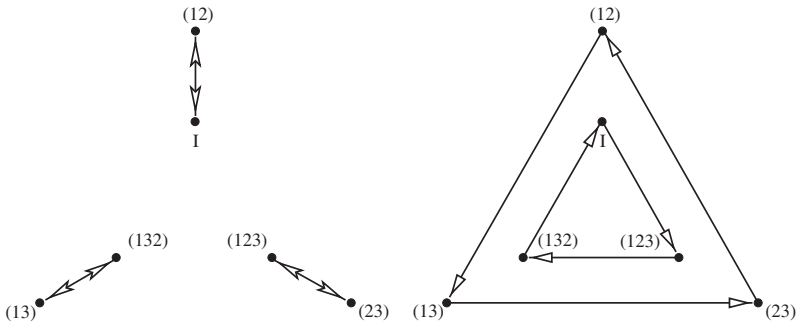


Fig. 1.1. The permutation of SYM_3 induced by (12) is shown on the left, and the permutation induced by (123) is shown on the right.

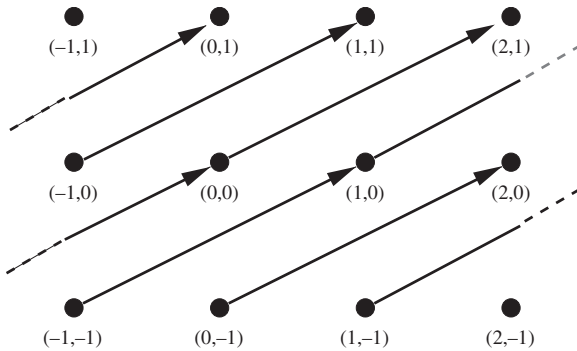


Fig. 1.2. The action of $(2, 1)$ on $\mathbb{Z} \oplus \mathbb{Z}$.

Example 1.8. In most introductions to group theory, Cayley's Basic Theorem is stated for *finite* groups. But we made no such assumption in our statement and the same proof as is given for finite groups works for infinite groups. Consider for example the direct product of two copies of the group of integers, $G = \mathbb{Z} \oplus \mathbb{Z}$. Here elements are represented by pairs of integers, and the binary operation is coordinatewise addition: $(a, b) + (c, d) = (a + c, b + d)$. In Figure 1.2 we have arranged the vertices corresponding to elements of G as the integral lattice in the plane. The arrows indicate the permutation of the elements of $\mathbb{Z} \oplus \mathbb{Z}$ induced by the element $(2, 1)$.

1.2 Graphs

One of the key insights into the study of groups is that they can be viewed as symmetry groups of graphs. We refer to this as “Cayley’s Better Theorem,” which we prove in Section 1.5.2. In this section we establish some terminology from graph theory, and in the following section we discuss groups acting on graphs.

Definition 1.9. A *graph* Γ consists of a set $V(\Gamma)$ of *vertices* and a set $E(\Gamma)$ of *edges*, each edge being associated to an unordered pair of vertices by a function “ENDS”: $\text{ENDS}(e) = \{v, w\}$ where $v, w \in V$. In this case we call v and w the *ends* of the edge e and we also say v and w are *adjacent*.

We allow the possibility that there are multiple edges with the same associated pair of vertices. Thus for two distinct edges e and e' it can be the case that $\text{ENDS}(e) = \text{ENDS}(e')$. We also allow loops, that is, edges whose associated vertices are the same. Graphs without loops or multiple edges are *simple* graphs.

Graphs are often visualized by making the vertices points on paper and edges arcs connecting the appropriate vertices. Two simple graphs are shown in Figure 1.3; a graph which is not simple is shown in Figure 1.4.

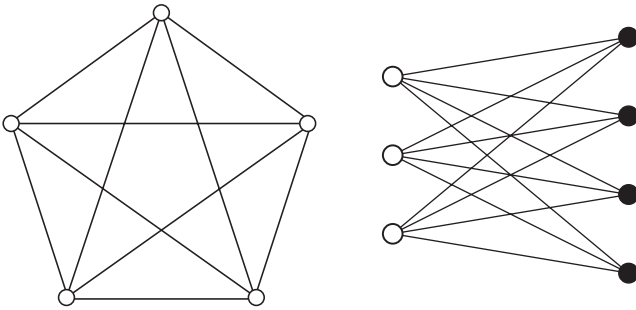


Fig. 1.3. The complete graph on five vertices, K_5 , and the complete bipartite graph $K_{3,4}$.

There are a number of families of graphs that arise in mathematics. The *complete graph* on n vertices has exactly one edge joining each pair of distinct vertices, and is denoted K_n . At the opposite extreme are the *null graphs*, which have no edges.

A graph is *bipartite* if its vertices can be partitioned into two subsets – by convention these subsets are referred to as the “black” and “white”

vertices – such that, for every $e \in E(\Gamma)$, $\text{ENDS}(e)$ contains one black vertex and one white vertex. The *complete bipartite* graphs are simple graphs whose vertex sets have been partitioned into two collections, V_\circ and V_\bullet , with edges joining each vertex in V_\circ with each vertex in V_\bullet . If $|V_\circ| = n$ and $|V_\bullet| = m$ then the corresponding complete bipartite graph is denoted $K_{n,m}$.

The *valence* or *degree* of a vertex is the number of edges that contain it. For example, the valence of any vertex in K_n is $n - 1$. If a vertex v is the vertex for a loop, that is an edge e where $\text{ENDS}(e) = \{v, v\}$, then this loop contributes twice to the computation of the valence of v . For example, the valence of the leftmost vertex in the graph shown in Figure 1.4 is six.

A graph is *locally finite* if each vertex is contained in a finite number of edges, that is, if the valence of every vertex is finite.

An *edge path*, or more simply a *path*, in a graph consists of an alternating sequence of vertices and edges, $\{v_0, e_1, v_1, \dots, v_{n-1}, e_n, v_n\}$ where $\text{ENDS}(e_i) = \{v_{i-1}, v_i\}$ (for each i). A graph is *connected* if any two vertices can be joined by an edge path. In Figure 1.4 we have indicated an

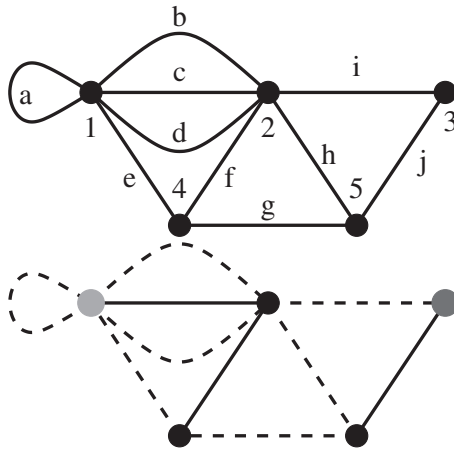


Fig. 1.4. On top is a graph which is not simple, with its vertices labelled by numbers and its edges labelled by letters. Below is the set of edges traversed in an edge path, joining the vertex labelled 1 to the vertex labelled 3, is indicated.

edge path from the leftmost vertex to the rightmost vertex. If v_i is the vertex labelled i and e_α is the edge labelled α , then this path is:

$$\{v_1, e_a, v_1, e_e, v_4, e_g, v_5, e_h, v_2, e_b, v_1, e_d, v_2, e_i, v_3\}$$

Cambridge University Press

978-0-521-89545-3 - Groups, Graphs and Trees: An Introduction to the Geometry of Infinite Groups

John Meier

Excerpt

[More information](#)

This notation is obviously a bit cumbersome and is only feasible when the graphs are small and the paths are short. If the graph is a simple graph, then one really only needs to list the sequence of vertices, which is a small economy of notation. In general, we will not need this level of specificity in dealing with edge paths in graphs.

A *backtrack* is a path of the form $\{v, e, w, e, v\}$ where one has travelled along the edge e and then immediately returned along e . A path is *reduced* if it contains no backtracks.

A *cycle* or *circuit* is a non-trivial edge path whose first and last vertices are the same, but no other vertex is repeated. The following paths in the graph shown in Figure 1.4 are all cycles:

1. $\{v_1, e_e, v_4, e_f, v_2, e_c, v_1\}$,
2. $\{v_1, e_b, v_2, e_d, v_1\}$,
3. $\{v_1, e_a, v_1\}$.

In Chapter 3 we study various groups that are closely connected to trees. A *tree* is a connected graph with no cycles. If you have not encountered trees in a previous course, working through the following exercise will help you gain some intuition for trees.

Exercise 1.10. Prove that the following conditions on a connected graph Γ are equivalent.

1. Γ is a tree.
2. Given any two vertices v and w in Γ , there is a unique reduced edge path from v to w .
3. For every edge $e \in E(\Gamma)$, removing e from Γ disconnects the graph. (Note: Removing e does not remove its associated vertices.)
4. If Γ is finite then $\#V(\Gamma) = \#E(\Gamma) + 1$.

While there are a number of interesting results about finite trees, in this book we shall be mainly interested in infinite trees. In particular, in Chapter 3 we explore groups that act on certain infinite, symmetric trees.

Definition 1.11 (regular and biregular trees). A *regular m -tree* is a tree where every vertex has fixed valence m . For a given value of m there is only one regular m -tree, which we denote \mathcal{T}_m . Notice that, since every vertex has valence m , the tree \mathcal{T}_m is infinite when $m \geq 2$.

A graph is *biregular* if it is bipartite, and all the vertices in one class have fixed valence m and all the vertices in the other class have fixed

Cambridge University Press

978-0-521-89545-3 - Groups, Graphs and Trees: An Introduction to the Geometry of Infinite Groups

John Meier

Excerpt

[More information](#)

valence n . Thus, for example, the complete bipartite graphs are all biregular. Given the valences, there is a unique biregular tree, which we denote by $\mathcal{T}_{m,n}$. You can see an example in Figure 1.5.

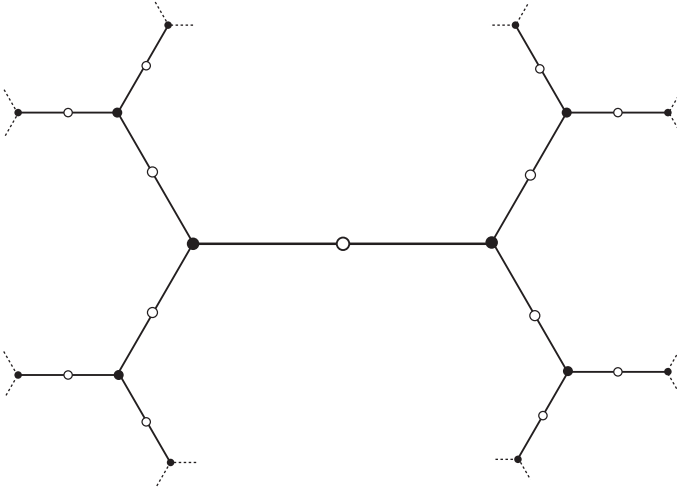


Fig. 1.5. The biregular tree $\mathcal{T}_{2,3}$.

It is often convenient to think of graphs as geometric objects where each edge is identified with the unit interval $[0, 1]$, where 0 corresponds to one of the associated vertices and 1 to the other. This convention allows us to refer to *midpoints* of edges, for example. The graphs shown in Figures 1.3, 1.4, and 1.5 have distorted this metric, meaning that if you measured the lengths of the edges in a given figure, those lengths will not all be the same. (This will occur in almost all of the graphs drawn in this book.) In addition to allowing us to specify points on edges, this geometric perspective allows us to think of paths as parametrized curves. Our convention will be that, given an edge path

$$\omega = \{v_0, e_1, v_1, e_2, \dots, e_n, v_n\}$$

in a graph Γ , there is an associated function $p_\omega : [0, 1] \rightarrow \Gamma$ where $p_\omega(i/n) = v_i$ and p_ω is linear when restricted to each of the subintervals $[i/n, (i+1)/n]$.

Remark 1.12. Our view of graphs is more topological than combinatorial. The reader who is a bit uneasy about our thinking of graphs as geometric objects, where edges have lengths, might want to look up the

Cambridge University Press

978-0-521-89545-3 - Groups, Graphs and Trees: An Introduction to the Geometry of Infinite Groups

John Meier

Excerpt

[More information](#)

10

Cayley's Theorems

definition of CW complexes, as we are viewing graphs as 1-dimensional CW complexes. We have not introduced this terminology or explicitly used the associated definition as it requires an understanding of topological spaces and quotient topologies.

There is one final variation on graphs that we will encounter in this text:

Definition 1.13. A *directed* graph consists of a vertex set V and an edge set E of ordered pairs of vertices. Thus each edge has an *initial* vertex and a *terminal* vertex. Graphically this direction is often indicated via an arrow on the edge. In thinking of directed graphs geometrically we assume that the initial vertex is identified with $0 \in [0, 1]$ and the terminal vertex with $1 \in [0, 1]$.

If we say a directed graph is connected we mean the underlying undirected graph is connected. (One can study “directed-connectedness” but that will not be relevant for us.)

In addition to directions on the edges, there are other sorts of decorations one can add to a graph. For example, one can have a set of labels \mathcal{L} and a function $\ell_V : V(\Gamma) \rightarrow \mathcal{L}$ that provides a labelling of the vertices. Or one could label the edges via $\ell_E : E(\Gamma) \rightarrow \mathcal{L}$, where the set of labels might be the same or different than the labels for the vertices. As an example, in Figure 1.4 we have shown a labelled graph where the vertices have been labelled with numbers and the edges with lower case letters.

1.3 Symmetry Groups of Graphs

Many important finite groups arise as symmetry groups of geometric objects. The dihedral groups are the symmetry groups of regular n -gons; the symmetric group SYM_n is isomorphic to the symmetry group of the regular $(n - 1)$ -dimensional simplex, for example the convex hull of

$$\{(1, 0, 0, \dots, 0), (0, 1, 0, \dots, 0), \dots, (0, \dots, 0, 1)\} \subset \mathbb{R}^n ;$$

the alternating subgroup A_n is isomorphic to the subgroup of symmetries of the regular $(n - 1)$ -dimensional simplex consisting of rotations in \mathbb{R}^n .

In this section we explore a similar theme, namely, we explore symmetry groups of graphs.

Definition 1.14. A *symmetry* of a graph Γ is a bijection α taking vertices to vertices and edges to edges such that if $\text{ENDS}(e) = \{v, w\}$, then