1

Introduction and Constructions

Let us begin with an informal example. By informal, we mean that we will use terms intuitively, as opposed to formally defining every term we use. Let us consider the symmetries of the cube. This is the usual cube or, if you prefer to have one in your hand, a six-sided die or a Rubik's cube. A common theme that we will see throughout this book, is that when considering symmetries of a graph, it is really helpful to have a clever labeling of the vertex set. As this is a book about symmetries of graphs, to us a clever labeling will be one that shows us, without too much work, many of the symmetries of the graph. We choose to label the vertices of the cube with elements of $\mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2$ as in Figure 1.1.

Notice that two vertices of the cube, as labeled in Figure 1.1, are adjacent if and only if their labels are different in exactly one coordinate. Also, for any face of the cube, there is a four-step rotation of the face in either the clockwise or counter-clockwise direction that is a symmetry of the cube (and of course the opposite face is rotated in the same fashion). As interchanging two opposite

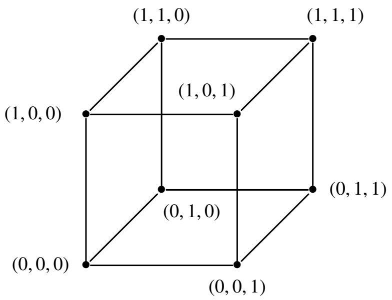

Figure 1.1 The 3-cube.

faces is also a symmetry of the cube, we may map any vertex of the cube to any other vertex of the cube. In the language we will introduce later, this is said as “the cube is vertex-transitive.”

In general, once one knows a particular graph is vertex-transitive, a next step in considering all of its symmetries is to think about distinguished *sets* of vertices. By “distinguished,” we mean that we can tell the sets are different somehow, under symmetry. We use the word “sets” here as vertex-transitive graphs have no distinguished vertices. For the cube, this means that we cannot distinguish (except by labeling) any corner of the cube from any other corner of the cube. But notice that for every vertex v of the cube, there is a unique vertex of distance three from v. Also, under any symmetry of a graph, vertices at some fixed distance in the graph are at the same distance in the graph under the symmetry. So these four pairs of vertices at distance three are permuted amongst themselves (these are the pairs $(0,0,0)$ and $(1,1,1)$, $(0,0,1)$ and $(1,1,0)$, $(0,1,0)$ and $(1,0,1)$, and $(1,0,0)$ and $(0,1,1)$. They are easy to remember as the vertex at distance three to (i,j,k) is simply $(i+1, j+1, k+1)$ with arithmetic in each coordinate performed modulo 2). These pairs of vertices are also called *antipodal* vertices. This term is borrowed from geography, where an antipodal point on the earth is the point on the earth opposite the given point. In the language of group theory, this says that the symmetries of the cube *act* on the pairs of antipodal points. We will later also say that a pair of antipodal vertices is a *block* of the symmetries of the cube, and the symmetries of the cube also permute the blocks.

A next obvious step might be to determine the symmetries of the antipodal vertices. That is, to determine how pairs of antipodal vertices are mapped to each other, and also if there are any symmetries, other than the identity, that fix each pair of antipodal points. For symmetries that fix each pair of antipodal points, consider the function that adds 1 modulo 2 in each coordinate of a vertex; that is $f\colon \mathbb{Z}_2^3 \to \mathbb{Z}_2^3$ given by $f(i,j,k) = (i+1, j+1, k+1)$. It is hopefully clear that if two vertices u and v are different in one coordinate, then $f(u)$ and $f(v)$ differ only in exactly the same coordinate. So f is a symmetry of the cube different from the identity that fixes each set of antipodal points. It turns out (we will prove this later), that there are no other such symmetries.

For symmetries of the antipodal vertices, there are four such pairs, so the largest number of such symmetries we can encounter is all of them! That is what we will see happens here. First, observe that no pair of antipodal points is contained in a face of the cube. So if we take a symmetry of a face, or square, and extend that to a symmetry of antipodal vertices (which can always be done as there is a unique antipodal vertex not on a face for each vertex on a face), we will have a symmetry of the cube. The symmetries of a face are just symmetries of the regular 4-gon, so four rotations and four reflections. This gives eight symmetries of the cube. To obtain elements of order 3, it is easiest to imagine that you are holding a cube between two fingers, with a finger on each antipodal vertex. Now spin the cube 120°! The cube does not change, so this gives a symmetry, and this can be repeated three times. We will see later that this is enough to obtain all symmetries of the cube, up to products of symmetries. That is, there are other symmetries, but these other symmetries can be obtained by successively applying symmetries we have already described. Cheating more than a little bit, the symmetries of the pairs of antipodal vertices is all symmetries of a set of 4 elements, so of order 24, while there are two symmetries that fix each set of antipodal vertices. In group-theoretic language, we have the group $\mathcal{S}_4 \times \mathbb{Z}_2$ of order 48.

We now turn to the main topics of this chapter. Section 1.1 is mainly concerned with the basic notation and terminology that we will use throughout the book. As the mathematical area that concerns symmetries is group theory, and more specifically permutation group theory, this will include terminology regarding both group theory and graph theory, as well as some examples and basic results.

The central goal of the rest of the chapter is to get examples of vertex-transitive graphs. An old result of Erdős and Rényi (1963) is that almost all graphs have a trivial automorphism group. Thus having any symmetry in a graph is rare, and so instead of looking for graphs with symmetries, we usually start with the symmetries, and then construct graphs that have those symmetries. The most common such construction is a Cayley graph, the topic of Section 1.2. We then give two ways of constructing every vertex-transitive digraph, namely double coset digraphs and orbital digraphs in Sections 1.3 and 1.4, respectively. While these techniques are formally different, as they both construct every vertex-transitive digraph, they should be thought of as being in some sense "the same." We then give the construction for the second most common family of vertex-transitive digraphs, the metacirculant digraphs, in Section 1.5.

One last editorial comment before we get going. You may have noticed we are a little confused in our use of the term "graph" and "digraph." Our main motivation in writing this book is to introduce you to symmetries of graphs. So why do we have digraphs? There are two simple answers. First, there are numerous results in the literature that deal with graphs only, where the proofs hold equally well for digraphs with such dramatic changes as replacing "graph" with "digraph." Hence, in many cases (but not all!), there really is no difference between graphs and digraphs, so why not work with the more general object? Second, there are also many times where a result is only proven for graphs, but in practice the result is needed for digraphs. The general philosophy is then that we are mainly interested in graphs, but when it comes time to prove a theorem, it is best to prove the theorem for digraphs if that is possible.

1.1 Basic Definitions

As this book is about groups acting on graphs, we begin with some of the basic definitions and results from permutation group theory and graph theory.

We emphasize that our permutation multiplication is on the left. That is, $fg(x) = f(g(x))$. Readers should be aware that sometimes in the literature permutation multiplication is written on the right.

A **permutation group** is a subgroup of the **symmetric group** on n letters, denoted $\mathcal{S}_n$. Unless otherwise stated, we will take the n letters that $\mathcal{S}_n$ permutes to be the elements of the set $\mathbb{Z}_n$, the integers modulo n. We denote the group of units in $\mathbb{Z}_n$ under multiplication by $\mathbb{Z}_n^*$, and note that $\mathrm{Aut}(\mathbb{Z}_n) = \{x \mapsto ax : a \in \mathbb{Z}_n^*\}$, the automorphism group of $\mathbb{Z}_n$. More generally, an **action** of a group G on a set X is a function $f\colon G \times X \to G$, with $f(g, x)$ written gx, such that $g(hx) = (gh)x$ and $1x = x$ for every $g, h \in G$ and $x \in X$ (of course 1 is the identity element in G). We will say that G **acts** on X on the left. In this text, unless otherwise stated, all groups and sets are finite, in which case the **degree** of an action is $|X|$, the number of elements in X. A related notion is that of a **permutation representation** of a group G, which is a homomorphism $\phi\colon G \to \mathcal{S}_n$ for some n. A standard result in group theory is that any action of G on X induces a homomorphism $\phi\colon G \to \mathcal{S}_X$ (Dummit and Foote, 2004, Proposition 4.1.1). So any action of G on X induces a corresponding permutation representation of G. We will occasionally abuse terminology and refer to $\phi(G)$ as a permutation representation if the action is clear. An action of G on X is called **faithful** if $\mathrm{Ker}(\phi) = 1$, where the **kernel of the action** ϕ of G on X is denoted $\mathrm{Ker}(\phi)$.

Example 1.1.1 Let n be a positive integer, and define $f\colon \mathbb{Z}_n \times \mathbb{Z}_n \to \mathbb{Z}_n$ by $f(g, x) = g + x$. For $g, h \in \mathbb{Z}_n$ and $x \in \mathbb{Z}_n$, $g + (h + x) = (g + h) + x$ and $0 + x = x$ so that f is an action of $\mathbb{Z}_n$ on itself. The degree of this action is $|\mathbb{Z}_n| = n$, and the action is faithful as if $f(g, x) = x$ for all $x \in \mathbb{Z}_n$, then $g = 0$. The corresponding permutation representation of $\mathbb{Z}_n$ is $\phi\colon \mathbb{Z}_n \to \mathcal{S}_n$ given by $\phi(g)$ is the function defined by $x \mapsto g + x \pmod n$. Note that $\phi(\mathbb{Z}_n)$ is the subgroup of $\mathcal{S}_n$ generated by the n-cycle $(0, 1, 2, \ldots, n-1)$, and is usually denoted $(\mathbb{Z}_n)_L$ (see Definition 1.2.12). Similarly, if $H \leq \mathbb{Z}_n$ is a subgroup of order m, then similar arguments show that $k\colon \mathbb{Z}_n \times (\mathbb{Z}_n/H) \to \mathbb{Z}_n \times (\mathbb{Z}_n/H)$ given by $k(g, x + H) = g + (x + H)$

is also an action. If $m > 1$ this action is not faithful, as $k(h, x + H) = x + H$ for every $h \in H$. The degree of this action is then n/m. The corresponding permutation representation of $\mathbb{Z}_n$ is $\delta\colon \mathbb{Z}_n \to \mathcal{S}_{n/m}$ given by $\delta(g)$ is the function defined by $x \mapsto x + g \pmod{n/m}$, and is also isomorphic to $(\mathbb{Z}_{n/m})_L$.

Definition 1.1.2 A permutation group $G \le \mathcal{S}_n$ is **transitive** if whenever $x, y \in \mathbb{Z}_n$, then there exists $g \in G$ such that $g(x) = y$.

Typically when discussing permutation groups, we will either start with a subgroup of $\mathcal{S}_n$, or will specify a group and an action of that group on a set X that then induces a natural subgroup of $\mathcal{S}_X$. Similarly, a concept about a permutation group translates into a concept about actions and vice versa, and we will usually refrain from defining analogous concepts in each context. So an action of G on X is **transitive** if for every $x, y \in X$ there is a $g \in G$ such that $gx = y$, and the **degree** of a transitive permutation group $G \le \mathcal{S}_n$ is n.

Example 1.1.3 Both of the actions in Example 1.1.1 are transitive as (using the notation from that example) if $x, y \in \mathbb{Z}_n$ then $f(y-x, x) = y$, and $k(y-x, x+H) = y + H$.

Definition 1.1.4 Let $G \le \mathcal{S}_n$ be transitive, and let the point $x \in \mathbb{Z}_n$. The **stabilizer of** x **in** G, denoted $\mathrm{Stab}_G(x)$, is defined by $\mathrm{Stab}_G(x) = \{g \in G : g(x) = x\}$. That is, $\mathrm{Stab}_G(x)$ is the set of all permutations in G that map x to x.

The stabilizer of x in G is often denoted G_x, and is a subgroup of G (Exercise 1.1.5).

Recall that every permutation in $\mathcal{S}_n$ can be written as a product of transpositions, and that this number is always even or always odd.

Definition 1.1.5 A permutation $\rho \in \mathcal{S}_n$ is **even** if it can be written as a product of an even number of transpositions, and **odd** if it can be written as a product of an odd number of transpositions. The set of all even permutations in $\mathcal{S}_n$ is a subgroup, called the **alternating group on** n **letters**, and is denoted A_n.

Also recall that a cycle of odd length is an even permutation, while a cycle of even length is an odd permutation. For proofs of the above mentioned facts and other information regarding the alternating group (see Dummit and Foote (2004, Section 3.5)).

Example 1.1.6 $\mathrm{Stab}_{\mathcal{S}_n}(n-1) = \mathcal{S}_{n-1}$, $\mathrm{Stab}_{A_n}(n-1) = A_{n-1}$, and $\mathrm{Stab}_{(\mathbb{Z}_n)_L}(x) = 1$, $x \in \mathbb{Z}_n$.

We now turn to some basic properties of the stabilizer.

Theorem 1.1.7 *Let $G \leq \mathcal{S}_n$, $x \in \mathbb{Z}_n$, and $h \in G$. Then $h\mathrm{Stab}_G(x)h^{-1} = \mathrm{Stab}_G(h(x))$. Consequently, if G is transitive, then $\mathrm{Stab}_G(x)$ is conjugate in G to $\mathrm{Stab}_G(y)$ for every $y \in \mathbb{Z}_n$.*

Proof Observe that

$$\begin{aligned}\mathrm{Stab}_G(h(x)) &= \{g \in G : g(h(x)) = h(x)\}\\ &= \{g \in G : h^{-1}gh(x) = x\}\\ &= \{g \in G : h^{-1}gh \in \mathrm{Stab}_G(x)\}\\ &= \{g \in G : g \in h\mathrm{Stab}_G(x)h^{-1}\}\\ &= h\,\mathrm{Stab}_G(x)\,h^{-1}.\end{aligned}$$

For the second statement, as G is transitive, there exists $h \in G$ such that $h(x) = y$. Then $\mathrm{Stab}_G(y) = \mathrm{Stab}_G(h(x)) = h\,\mathrm{Stab}_G(x)\,h^{-1}$. □

The following result is quite useful, and is sometimes called the orbit-stabilizer theorem.

Theorem 1.1.8 *Let $G \leq \mathcal{S}_n$, and $x \in \mathbb{Z}_n$. The set $G(x) = \{g(x) : g \in G\}$, is the* **orbit of x in G**. *Then $|G| = |G(x)| \cdot |\mathrm{Stab}_G(x)|$, or equivalently, $|G(x)| = [G : \mathrm{Stab}_G(x)]$.*

Proof Define $\phi\colon G \to G(x)$ by $\phi(g) = g(x)$. Note that ϕ is onto (or surjective) by definition, and so $|\phi(G)| = |G(x)|$. Also,

$$\begin{aligned}\phi(g) = \phi(h) &\text{ if and only if } g(x) = h(x)\\ &\text{ if and only if } h^{-1}g(x) = x\\ &\text{ if and only if } h^{-1}g \in \mathrm{Stab}_G(x)\\ &\text{ if and only if } h^{-1}g\mathrm{Stab}_G(x) = 1 \cdot \mathrm{Stab}_G(x) \text{ (as left cosets)}\\ &\text{ if and only if } h \text{ and } g \text{ are in the same left coset of } \mathrm{Stab}_G(x).\end{aligned}$$

Thus $|\phi(G)|$ is the number of left cosets of $\mathrm{Stab}_G(x)$ in G, and so $|G(x)| = [G : \mathrm{Stab}_G(x)]$. As all groups here are finite, $[G : \mathrm{Stab}_G(x)] = |G|/|\mathrm{Stab}_G(x)|$. □

The following application of the orbit-stabilizer theorem is originally due to Miller (1903).

Theorem 1.1.9 *Let G be a transitive group of degree n, and p a prime. The highest power p^k of p dividing n also divides the length of every orbit of a Sylow p-subgroup of G.*

Proof Let P be a Sylow p-subgroup of G, and $x \in \mathbb{Z}_n$. By the orbit-stabilizer theorem, $n = [G : \mathrm{Stab}_G(x)]$ and $[P : \mathrm{Stab}_P(x)] = |P(x)|$. Then,

$$
\begin{aligned}
n \cdot [\mathrm{Stab}_G(x) : \mathrm{Stab}_P(x)] &= [G : \mathrm{Stab}_G(x)] \cdot [\mathrm{Stab}_G(x) : \mathrm{Stab}_P(x)] \\
&= [G : \mathrm{Stab}_P(x)] \\
&= [G : P] \cdot [P : \mathrm{Stab}_P(x)] \\
&= [G : P] \cdot |P(x)|.
\end{aligned}
$$

As $\gcd([G : P], p) = 1$, where $\gcd(m, n)$ is the greatest common divisor of m and n, and p^k divides n, we see that p^k divides $|P(x)|$. □

The following immediate corollary is the form of the previous result used most often (for our purposes), and is also due to Miller.

Corollary 1.1.10 *Let p be prime and $k \geq 1$. A Sylow p-subgroup of a transitive group of degree p^k is transitive.*

Definition 1.1.11 A permutation group $G \leq \mathcal{S}_n$ is **semiregular** if $\mathrm{Stab}_G(x) = 1$ for every $x \in \mathbb{Z}_n$, and G is **regular** if G is both semiregular and transitive.

The next result follows directly from the orbit-stabilizer theorem as if $G \leq \mathcal{S}_n$ is transitive, then G only has one orbit.

Corollary 1.1.12 *A transitive group is regular if and only if its order and degree are the same.*

Example 1.1.13 The group $\mathcal{S}_n$ is regular if and only if $n \leq 2$, and A_n is regular if and only if $n = 1$ or 3 but is semiregular for $n = 2$. The group $(\mathbb{Z}_n)_L$ is regular for all positive integers n.

Solution The group $\mathcal{S}_n$ is regular by Corollary 1.1.12 if and only if $n! = n$, which is true if and only if $n \leq 2$. Also, $A_1 = 1$ is regular, $A_2 = 1$ is not regular, and for $n \geq 3$, the group A_n is transitive and so regular if and only if $n!/2 = n$, which is true if and only if $n = 3$. Finally, $(\mathbb{Z}_n)_L$ is transitive of order n and so regular. □

A **digraph** Γ is an ordered pair (V, A), where V is a nonempty set (V is the **vertex set of** Γ), and $A \subseteq \{(u, v) : u, v \in V\}$ of ordered pairs of V (A is the **arc set of** Γ). An arc (u, v) of a digraph is usually thought of as being directed from u to v. If $(u, v) \in A(\Gamma)$ and $(v, u) \in A(\Gamma)$, we can identify these two arcs and consider it an **edge** (a loop is also considered an edge). A **graph** Γ is a digraph in which $(u, v) \in A(\Gamma)$ if and only if $(v, u) \in A(\Gamma)$, and in this case we think of the edges as being unordered pairs of vertices, denoted uv. The arc set of a graph Γ is usually called the **edge set** of Γ and denoted by $E(\Gamma)$ instead of $A(\Gamma)$. According to this definition, multiple edges (different edges with the same endpoints) are not allowed. This is not done from any sort of dislike of multiple edges, but more from a desire to start with the simplest definition of a graph, and in most, but not all, situations concerning symmetry in digraphs, multiple edges are irrelevant. This definition does allow loops, but similarly in this book loops are irrelevant most of the time. There are many situations, some of which will be encountered later, where the context calls for modifications to the definition of a graph. In sections or chapters where multiple edges are needed, this will be stated explicitly at the beginning of the section or chapter. In this text, all graphs are finite (so have finite vertex sets). Basic digraph and graph terms and operations such as walks, intersection of digraphs, etc., are defined as usual. See Bollobás (1998), for example.

As is customary, for a graph Γ we denote by $N_\Gamma(u) = \{v : uv \in E(\Gamma)\}$ the set of neighbors in Γ of the vertex u. For a digraph Γ, we denote the outneighbors of u by $N^+_\Gamma(u)$, and the inneighbors by $N^-_\Gamma(u)$. That is, $N^+_\Gamma(u) = \{v : (u,v) \in A(\Gamma)\}$ and $N^-_\Gamma(u) = \{v : (v,u) \in A(\Gamma)\}$. If the graph or digraph Γ is clear, we may simply write $N^+(u)$, etc.

Definition 1.1.14 An **isomorphism** between two digraphs Γ_1 and Γ_2, is a bijection $\phi\colon V(\Gamma_1) \to V(\Gamma_2)$ such that $(\phi(u), \phi(v)) \in A(\Gamma_2)$ if and only if $(u, v) \in A(\Gamma_1)$. We write $\Gamma_1 \cong \Gamma_2$.

Thus an isomorphism is a one-to-one mapping of the vertex set onto the vertex set that preserves arcs.

Any isomorphism ϕ between two digraphs Γ_1 and Γ_2 induces a natural bijection between $A(\Gamma_1)$ and $A(\Gamma_2)$ given by $(u, v) \mapsto (\phi(u), \phi(v))$. We will often abuse notation and simply write $\phi(u, v)$ instead of $(\phi(u), \phi(v))$ for the image of the arc (u, v) under ϕ.

Definition 1.1.15 Any bijection ϕ from $V(\Gamma_1)$ to a set X induces a digraph on X isomorphic to Γ_1. Namely, define a digraph Γ_2 by $V(\Gamma_2) = X$ and $A(\Gamma_2) = \{(\phi(u), \phi(v)) : (u, v) \in A(\Gamma_1)\}$. We will adopt the notational convention of setting $\phi(\Gamma_1) = \Gamma_2$.

An **automorphism** of a digraph is an isomorphism of a digraph with itself. The set of all automorphisms of a digraph Γ is a group under function composition (or permutation multiplication if you prefer – see Exercise 1.1.8), and is called the **automorphism group of** Γ, denoted $\mathrm{Aut}(\Gamma)$. Of course, $\mathrm{Aut}(\Gamma)$ also acts on the arcs of a digraph and the edges of a graph.

Definition 1.1.16 A digraph whose automorphism group is transitive on its vertex set is called a **vertex-transitive digraph**, and a digraph whose automorphism group is transitive on its arc set is called an **arc-transitive digraph**. A graph Γ is **edge transitive** if $\mathrm{Aut}(\Gamma)$ is transitive on its edge set $E(\Gamma)$.

Example 1.1.17 A complete graph K_n of order $n \geq 1$ is vertex-transitive, arc-transitive, and $\text{Aut}(\Gamma) = \mathcal{S}_n$.

Solution As K_n has every arc, if $\sigma \in \mathcal{S}_n$ then $\sigma(u, v) \in A(K_n)$ for every pair of distinct vertices u and v. Thus $\text{Aut}(K_n) = \mathcal{S}_n$ (as $\mathcal{S}_n$ is the "largest" permutation group on n vertices). The group $\mathcal{S}_n$ is certainly transitive, and also if $(u_1, v_1), (u_2, v_2) \in A(\mathcal{S}_n)$, then there is $\sigma \in \mathcal{S}_n$ with $\sigma(u_1, v_1) = (u_2, v_2)$. Thus K_n is vertex-transitive and arc-transitive. □

Example 1.1.18 For $n \geq 1$, define a graph Q_n by $V(Q_n) = \mathbb{Z}_2^n$ and two vertices are adjacent if and only if they differ in exactly one coordinate. The graph Q_n is the ***n*-dimensional hypercube** or the ***n*-cube** and is vertex-transitive. The 3-cube, whose automorphism group was discussed at the beginning of this chapter, is shown in Figure 1.1.

Solution Define $\tau_i \colon \mathbb{Z}_2^n \to \mathbb{Z}_2^n$ by $\tau_i(x)$ adds 1 modulo 2 in the ith coordinate of x and is the identity on the other coordinates. If $e \in E(Q_n)$, then $\tau_i(e) \in E(Q_n)$, as in any coordinate, adding a constant (either 0 or 1) does not change whether the coordinates are the same or different. Thus $\tau_i \in \text{Aut}(Q_n)$ for every $1 \leq i \leq n$. Additionally, it is easy to see that any element of $\mathbb{Z}_2^n$ can be transformed into any other element of $\mathbb{Z}_2^n$ by changing (or adding 1) coordinates of the first that differ with the second. Thus $\langle \tau_i : 1 \leq i \leq n \rangle$, the subgroup of $\text{Aut}(Q_n)$ generated by the τ_i, is transitive and Q_n is vertex-transitive. □

The $(n + 1)$-cube can be constructed from the n-cube by taking two copies of the n-cube and joining them by a 1-factor. The vertices of one n-cube has an additional coordinate added with a 1 in the additional coordinate, while the vertices of the other n-cube has an additional coordinate added with a 0 in the additional coordinate. The edges of the 1-factor are then between corresponding vertices in the two copies. That is, the edges of the 1-factor only differ in the additional coordinate. This can be seen in Figure 1.1, where the additional coordinate that is added is the first coordinate, and the two copies of the 2-cube are the top and bottom faces of the cube. This construction for the $(n + 1)$-cube from the n cube is often useful for induction arguments.

Example 1.1.19 The Petersen graph given in Figure 1.2 is vertex-transitive.

Solution It is easy to see that a 72° rotation leaves the Petersen graph invariant, and so any vertex of the "outside" 5-cycle $0, 1, 2, 3, 4, 0$ can be mapped to any other vertex on the "outside" 5-cycle and similarly, any vertex of the "inside" 5-cycle $0', 1', 2', 3', 4', 0'$ can be mapped to any other vertex on the "inside" 5-cycle. It thus only remains to show that there is an

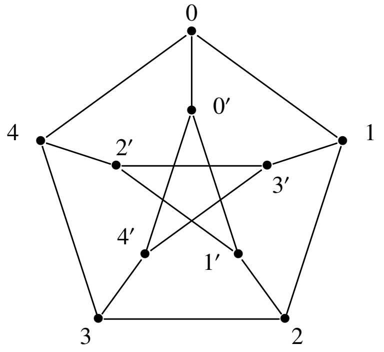

Figure 1.2 The Petersen graph.

automorphism of the Petersen graph that interchanges the outside 5-cycle and the inside 5-cycle. Consider the permutation $(0,0')(1,1')(2,2')(3,3')(4,4')$. Straightforward computations (which you should do!) verify that this map is an automorphism of the Petersen graph and so the Petersen graph is vertex-transitive. □

A word about our notation for cycles. Usually a cycle is denoted by a sequence of vertices, say $v_0v_1,\ldots,v_{n-1}v_0$. You will notice in the previous paragraph we used the slightly nonstandard notation of $v_0,v_1,\ldots,v_{n-1},v_0$, as when using numbers for the labels of vertices, there can often be ambiguity – particularly if the graph has many vertices. For example, in a graph with vertex set $\mathbb{Z}_{25}$ the cycle 01230 could be a cycle of length 4 or two different cycles of length 3! So we will feel free to use commas between vertices in denoting a cycle, or not if it causes no ambiguity. A similar comment holds for cycle notation of permutations.

Before our next discussion, we will need some additional definitions.

Definition 1.1.20 Let Γ be a regular graph. The **valency of** Γ is the number of edges incident with any vertex. We say Γ is **cubic** if it is regular of valency 3. Finally, Γ is **symmetric** if its automorphism group is transitive in its action on both its vertex set and its arc set.

While in this text we are usually concerned with vertex-transitive graphs, from the very beginning of the study of vertex-transitive graphs there has been interest in graphs whose automorphism group also acts transitively on arcs or other subdigraphs in graphs. Indeed, starting in 1932, Foster started compiling what is now known as the **Foster census** (Foster, 1988) of cubic symmetric