1
Introduction

1.1 Construction and decomposition

Mathematical thinking has an interesting duality reflected in its fundamental techniques: constructing and decomposing. Whereas construction methods provide a means for producing larger objects from smaller ones, decomposition theorems enable us to identify the basic or irreducible blocks from which general mathematical structures can be obtained. A combination of these two methods usually leads to an understanding of how complicated objects are constructed from their building blocks. Perhaps the simplest example is in the category of sets. In this category we can construct a new set from given ones by taking their co-product, that is, their disjoint union. We can also easily decompose any given set as a co-product of smaller sets by identifying a partition. Similarly, in a first course on group theory, we teach the (external) direct product construction along with the criteria for identifying an (internal) direct product decomposition.

These algebraic and set theoretic examples are brought together in the study of group actions. A general construction for permutation groups takes two groups G and H acting on disjoint sets Γ and Δ and produces a permutation group isomorphic to $G \times H$ acting on the union $\Gamma \cup \Delta$. Using a more complicated method we may construct the wreath product $G \wr S_k$ acting on the disjoint union of k copies of the set Γ (see Section 5.2.1). Conversely, a first analysis of a permutation group G on a set Γ identifies the partition of Γ given by the G-orbits, and determines G as a subgroup of the direct product of the transitive permutation groups induced on these G-orbits (see Section 2.1). In this way, the class of permutation groups can be understood by focusing on the smaller class of transitive groups, and the subgroups of their direct products. In turn, many questions concerning transitive groups can be answered via a reduction to primitive groups, that is, to groups for which the point set has no non-trivial invariant co-product decomposition. Thus the co-product object of sets (that is, partitions of sets) serves as a very useful tool in the theory of permutation groups.

Equally important in the category of sets is the product object, which is the cartesian product. As above, there is a permutation group construction that, given groups G and H acting on Γ and Δ, respectively, returns a group isomorphic to $G \times H$ acting on $\Gamma \times \Delta$. Moreover, the wreath product $G \wr S_k$ also acts on Γ^k in its product action (see Section 5.2.2). This product action plays a very important role in the theory of primitive groups. For example, one interpretation of the O'Nan–Scott Theorem for finite primitive permutation groups is that there are four fundamental types of finite primitive groups and all others arise as subgroups of wreath products $G \wr S_k$ in product action where G is a fundamental primitive group. These four fundamental types also occur as maximal subgroups in finite symmetric and alternating groups; see Theorem 7.11. The O'Nan–Scott Theorem (see Chapter 7) has provided the most useful modern method for identifying the possible structures of finite primitive groups and is now used routinely for their analysis. Thus a cartesian decomposition concept, complementing the cartesian product construction, should play an important role in the study of permutation groups, especially in that of the finite primitive ones. It is therefore surprising that there is no widely used such cartesian decomposition concept. Instead, mathematicians usually work their way around introducing one, and this can lead to imprecise and inadequate treatment of groups acting in product action.

In our work we show how a properly defined cartesian decomposition concept leads to a new way of analysing (not necessarily finite) transitive permutation groups by decomposing them with respect to invariant cartesian decompositions of the point set. This leads to a new theory for the class of permutation groups with a transitive minimal normal subgroup. This theory is particularly powerful in the case when the minimal normal subgroup is of the form T^k for a simple group T and a positive integer k; in particular, this smaller class contains all finite primitive groups. The cartesian decomposition concept we use in this book first appeared in the paper by L. G. Kovács (Kovács 1989b) and was used to identify wreath decompositions of permutation groups. Kovács used the name *system of product-imprimitivity* for this concept, but we find the name *cartesian decomposition* more descriptive. Cartesian decompositions are exploited to deepen the current understanding of permutation groups contained in wreath products in product action. We demonstrate, in the finite case, that the study of cartesian decompositions combined with certain facts about finite simple groups, which depend on the finite simple group classification, leads to an unexpectedly detailed description of permutation groups that act on cartesian products. In particular, this description sheds new light on the theory of primitive permutation groups and also on the larger families of quasiprimitive and innately transitive permutation groups introduced in Section 1.3. It also applies, for example, to infinite primitive and quasiprimitive permutation groups with finite stabilisers; see Theorems 3.18 and 7.9.

1.2 Cartesian decompositions

To introduce the intuitive idea behind cartesian decompositions, consider the following example. In the three-dimensional coordinate system, the set C of points

$$\{(x_1, x_2, x_3) \mid x_1,\ x_2,\ x_3 \in \{0,1\}\}$$

forms a cube. The cube C lives in a 3-dimensional space and it can be bisected in three different ways using planes parallel to the three fundamental planes of the coordinate system. For example, four of the 8 points lie on the plane defined by the equation $x_1 = 0$ and four lie on the plane with equation $x_1 = 1$, and this gives the first partition. Similarly, the second partition is determined by the planes with equations $x_2 = 0$ and $x_2 = 1$, and the third partition is determined by the equations $x_3 = 0$ and $x_3 = 1$. Let us denote these partitions by Γ_1, Γ_2, and Γ_3, respectively. Then

$$\begin{aligned}
\Gamma_1 &= \{\{(0,0,0),(0,0,1),(0,1,0),(0,1,1)\},\\
&\quad\ \{(1,0,0),(1,0,1),(1,1,0),(1,1,1)\}\};\\
\Gamma_2 &= \{\{(0,0,0),(0,0,1),(1,0,0),(1,0,1)\},\\
&\quad\ \{(0,1,0),(0,1,1),(1,1,0),(1,1,1)\}\};\\
\Gamma_3 &= \{\{(0,0,0),(0,1,0),(1,0,0),(1,1,0)\},\\
&\quad\ \{(0,0,1),(0,1,1),(1,0,1),(1,1,1)\}\}.
\end{aligned}$$

In other words, the partition Γ_i divides the vertices of the cube C into two blocks according to the i-th coordinate.

We note that the first coordinate of a point is determined by its block in the first partition, and similarly its second and third coordinates are determined by its blocks in the second and the third partitions, respectively. Since each point is determined by its three coordinates we see that the intersection of three blocks, one from each of the three different Γ_i, has size one. In other words,

$$|\gamma_1 \cap \gamma_2 \cap \gamma_3| = 1 \quad \text{whenever} \quad \gamma_1 \in \Gamma_1,\ \gamma_2 \in \Gamma_2,\ \gamma_3 \in \Gamma_3.$$

This motivates the definition of a cartesian decomposition.

Definition 1.1 A *cartesian decomposition*, $\mathcal{E}$, of a set Ω is a finite set of partitions, $\mathcal{E} = \{\Gamma_1, \ldots, \Gamma_\ell\}$, of Ω such that $|\Gamma_i| \geqslant 2$ for each i and

$$|\gamma_1 \cap \cdots \cap \gamma_\ell| = 1 \quad \text{for each} \quad \gamma_1 \in \Gamma_1, \ldots, \gamma_\ell \in \Gamma_\ell.$$

A cartesian decomposition is said to be *trivial* if it contains only one partition, namely the partition into singletons. A cartesian decomposition is said to be *homogeneous* if all the Γ_i have the same cardinality.

If $\{\Gamma_1, \ldots, \Gamma_\ell\}$ is a cartesian decomposition of a set Ω, then the defining property yields a well defined bijection between Ω and $\Gamma_1 \times \cdots \times \Gamma_\ell$, given by

$$\omega \mapsto (\gamma_1, \ldots, \gamma_\ell) \tag{1.1}$$

where, for $i = 1, \ldots, \ell$, the block $\gamma_i \in \Gamma_i$ is the unique block of Γ_i which contains ω. Thus the set Ω can be naturally identified with the cartesian product $\Gamma_1 \times \cdots \times \Gamma_\ell$.

Let us now turn back to the example given before Definition 1.1, and view the cube C as a graph in which two vertices are joined if and only if they only differ in one position. Let W denote the wreath product of the cyclic group of order 2 and the symmetric group of degree 3. That is, $W = C_2 \wr \mathsf{S}_3 = (C_2 \times C_2 \times C_2) \rtimes \mathsf{S}_3$ and W is a group of order 48. The group W acts on C as follows. Set $B = C_2 \times C_2 \times C_2$ and let $x = (x_1, x_2, x_3) \in C$. The first copy of C_2 in B flips the first coordinate of x (that is, interchanges 0 and 1), the second copy of C_2 flips the second coordinate, while the third copy flips the third coordinate. The group S_3 permutes the coordinates of x naturally. Easy consideration shows that each element of W is an automorphism of this 'cube graph', and simple graph theoretic consideration shows that W is the full automorphism group of C. This example illustrates the strong relationship between wreath products and cartesian decompositions. In fact, the observations we made in this simple example are generalised in Section 12.2 where we will study the general relationship between cartesian products of graphs and cartesian decompositions.

1.3 Cartesian factorisations

As mentioned in Section 1.1, modern studies of finite primitive permutation groups identified groups preserving cartesian decompositions as having fundamental significance for a theory of primitive groups, especially in the O'Nan–Scott Theorem. They are of similar importance in studying larger families of permutation groups such as quasiprimitive groups and innately transitive groups. A permutation group is said to be *quasiprimitive* if all its non-trivial normal subgroups are transitive and it is called *innately transitive* if it has a transitive minimal normal subgroup. Hence every primitive group is quasiprimitive (Corollary 2.21), and every finite quasiprimitive group is innately transitive. Moreover, for a finite group, a minimal normal subgroup is a direct product of finitely many copies of a simple group (see Lemma 3.14). We consider finite and infinite innately transitive groups with a minimal normal subgroup of this kind. The purpose of this book is to present a theory of cartesian decompositions that are invariant under the action of such a group, and to use it to present characterisations of the primitive, quasiprimitive, and innately transitive groups having a minimal normal subgroup of this form. We will also apply this theory in various group theoretic and combinatorial contexts.

A central problem we wish to solve is the following.

> *For a given innately transitive group, decide if the group action can be realised on a non-trivial cartesian product of smaller sets; that is, decide if the group leaves invariant a non-trivial cartesian decomposition.*

Our approach involves a mixture of combinatorial and group theoretic methods. The example of the cube in the previous section shows that cartesian decompositions sometimes arise naturally. However, they are not always so easy to recognise, and the following example shows that the structure of the acting group may help in finding invariant cartesian decompositions of the underlying set. We use the primitive subgroup $G = \mathsf{Aut}(\mathsf{A}_6) \cong \mathsf{P\Gamma L}_2(9)$ of S_{36} for illustration. The socle T of this group is isomorphic to A_6. The reason why G acts on a cartesian product $\Gamma \times \Gamma$, with $|\Gamma| = 6$, is that T has two subgroups A and B, both isomorphic to A_5, such that $A \cap B$ is a point stabiliser T_ω, the subgroup G_ω swaps A and B, and T can be factorised as $T = AB$. Since invariant partitions of the point set correspond to overgroups of the point stabiliser (Lemma 2.14), the socle T preserves two partitions Γ_1, Γ_2 of the point set that are orthogonal in the sense that $\gamma_1 \cap \gamma_2$ is a singleton for all $\gamma_1 \in \Gamma_1$ and $\gamma_2 \in \Gamma_2$. In other words, the underlying set of G can be naturally identified with the cartesian product $\Gamma_1 \times \Gamma_2$. Further, as G_ω swaps A and B, G_ω also swaps these partitions. Hence the group $G = TG_\omega$ can also be viewed as a permutation group acting on the cartesian product $\Gamma_1 \times \Gamma_2$.

This example suggests that we may be able to solve our original problem by studying certain partitions of the point set invariant under a suitable minimal normal subgroup, and, in turn, such partitions may be pinpointed by understanding the factorisations of this minimal normal subgroup.

In this book, after giving a thorough treatment of the fundamental theory of permutation groups (of arbitrary cardinality), we focus on permutation groups G with a minimal normal subgroup M that is transitive on the underlying point set Ω. In other words, we focus on innately transitive permutation groups. Such a subgroup M is called a plinth of G, and is characteristically simple. One of our more fundamental results shows that each partition in a G-invariant cartesian decomposition of Ω is M-invariant (Theorem 8.3). In order to find M-invariant partitions of the point set Ω that form a G-invariant cartesian decomposition, we study the overgroups of a fixed point stabiliser in M. We show in Theorem 8.2 that if $\{\Gamma_1, \ldots, \Gamma_\ell\}$ is a G-invariant cartesian decomposition of Ω, $\gamma_1 \in \Gamma_1, \ldots, \gamma_\ell \in \Gamma_\ell$, and $\{\omega\} = \gamma_1 \cap \cdots \cap \gamma_\ell$, then

$$\bigcap_{i=1}^{\ell} M_{\gamma_i} = M_\omega \text{ and } M_{\gamma_i}\left(\bigcap_{j \neq i} M_{\gamma_j}\right) = M \text{ for all } i \in \{1, \ldots, \ell\}. \quad (1.2)$$

The collection of subgroups M_{γ_j} satisfying (1.2) is called a *cartesian factorisation*† of M, and is studied in Chapter 8. The set $\{A, B\}$ of subgroups of $T = \mathsf{A}_6$ identified in the example presented above is a

† Cartesian factorisations were referred to as *cartesian systems of subgroups* in our earlier papers (Baddeley, Praeger & Schneider 2004a, Baddeley, Praeger & Schneider 2004b, Baddeley, Praeger & Schneider 2006, Baddeley, Praeger & Schneider 2007, Praeger & Schneider 2007, Baddeley, Praeger & Schneider 2008, Praeger & Schneider 2012).

cartesian factorisation for A_6. Looking at the last displayed equation, it is clear that in order to find G-invariant cartesian decompositions of Ω, we need to study factorisations of the characteristically simple group M. This study is carried out under the stronger condition that M is a direct product $M = T_1 \times \cdots \times T_k$ where the T_i are simple groups. Hence we introduce the following definition.

Definition 1.2 A group M is said to be FCR (finitely completely reducible) if M can be written as a direct product $M = T_1 \times \cdots \times T_k$ where the T_i are simple groups. We sometimes say that M is an FCR-*group*.

Note that a finite characteristically simple group is FCR (Lemma 3.14) and so is a minimal normal subgroup of an infinite quasiprimitive permutation group with finite stabilisers (Theorem 3.18). In a finite FCR-group, the factors T_i are finite simple groups, and the machinery provided by the finite simple group classification is available for our use. In particular, using the available knowledge on factorisations of finite simple groups (Liebeck, Praeger & Saxl 1990, Baddeley & Praeger 1998), the factorisations occurring in relation to cartesian decompositions of finite sets can be characterised. Moreover, using this characterisation, in the most interesting cases, the G-invariant cartesian decompositions can be described (see Theorems 8.17, 9.7, 10.13).

1.4 Primitive, quasiprimitive and innately transitive groups: 'O'Nan–Scott theories'

One of the most important outcomes of studying invariant cartesian decompositions is a better understanding of the O'Nan–Scott theory of primitive, quasiprimitive and innately transitive groups. The *primitive* permutation groups on a set Ω are those which leave invariant only the trivial partitions of Ω: the partition in which each part consists of a single point, and the partition with just one part. Since the early 1980s the study of finite primitive permutation groups has been transformed by the O'Nan–Scott Theorem which identifies several types of finite primitive groups, and asserts that each finite primitive group is of one of these types. Proofs of the O'Nan–Scott theorem for finite primitive permutation groups can be found in (Scott 1980, Aschbacher & Scott 1985, Kovács 1986, Buekenhout 1988, Liebeck, Praeger & Saxl 1988) and more detailed treatments of it in (Dixon & Mortimer 1996, Cameron 1999).

Cameron's approach (Cameron 1999) strongly influenced the exposition in this book. He divides the finite primitive groups into two families: the groups in the first family are called the *basic groups*, and the other family is formed by the *non-basic primitive groups*. In Cameron's terminology, a basic group is one that cannot be embedded into a wreath product. This definition of basic groups led to a slight inaccuracy in his treatment of primitive groups. Namely, Cameron treats almost simple groups as basic, even though $\mathsf{P\Gamma L}_2(9)$, for instance, acting primitively on 36 points, as considered in Section 1.3, preserves a non-trivial homogeneous cartesian decomposition, and hence it can be embedded into a wreath product in product action. Nevertheless, it is true, and a useful fact, that each finite primitive permutation group is either basic, or a subgroup of a wreath product $H \wr S_k$ in product action on Γ^k, where H is a basic primitive group on Γ. The latter situation is equivalent to the existence of a non-trivial homogeneous cartesian decomposition preserved by the group. In this book we focus on this situation, but we extend our scope to the class of (possibly infinite) groups that have a transitive minimal normal subgroup.

Much recent work on finite primitive permutation groups concentrated on understanding the basic groups, especially those related to non-abelian simple groups and to irreducible representations of finite groups. This information together with the wreath product construction leads to the solution of many problems in algebra, number theory and combinatorics. However, for some applications, detailed information is needed on precisely which subgroups of a wreath product $H \wr S_k$ with H primitive on Γ, are themselves primitive on Γ^k. The seminal paper of Kovács (Kovács 1989a) introduced the concept of a 'blow-up' of a primitive group and provided criteria for identifying such subgroups for almost all types of primitive groups H. Moreover, this led, in 1990, to a classification (Praeger 1990) of all embeddings of finite primitive groups into wreath products in product action; that is, a classification of all homogeneous cartesian decompositions invariant under primitive groups. In his study of the finite lattice representation problem (see (Pálfy & Pudlák 1980)), Aschbacher (Aschbacher 2009a, Aschbacher 2009b) addressed the same questions and obtained a similar solution.

We extend Kovács's blow-up concept to a larger class of permutation groups in Section 11.1.

The product action of a wreath product was also a pivotal concept in describing finite quasiprimitive groups, a strictly larger class of groups than that of the primitive groups, and one which arises naturally in many combinatorial applications. The term was coined in the 1970s by

Wielandt (private communication from W. Knapp to the first author in January 1994) and first appeared in print in works of Knapp (for example in (Knapp 1973)). In fact Peter Neumann's critical analysis of the 2$^{\text{eme}}$ Mémoire of Évariste Galois, in (Neumann 2006), suggests that the 'primitive' permutation groups Galois studied in the early 19$^{\text{th}}$ century were in fact the quasiprimitive groups.

To date a great many applications of group theory in combinatorics and other subjects have depended upon a reduction to a case involving a primitive permutation group. However such a reduction is not always possible. The first author consequently initiated an investigation into the suitability of quasiprimitivity, rather than primitivity, as a reduction tool in applications of group theory. This investigation was two-pronged; it involved, on the one hand, an attempt to understand quasiprimitive permutation groups in a purely group-theoretic setting, and on the other, it also involved several applications of quasiprimitivity to classification problems for combinatorial structures, such as incidence geometries (Cara, Devillers, Giudici & Praeger 2012), line-transitive linear spaces (Camina & Praeger 2001), k-arc transitive graphs (Ivanov & Praeger 1993, Baddeley 1993b, Fang 1995, Li 2001, Hassani, Nochefranca & Praeger 1999, Praeger 1993), k-arc-transitive Cayley graphs (Li 2005), locally primitive graphs (Praeger, Pyber, Spiga & Szabó 2012), locally quasiprimitive graphs (Li, Praeger, Venkatesh & Zhou 2002), and strongly regular edge-transitive graphs (Morris, Praeger & Spiga 2009).

Understanding finite quasiprimitive groups in sufficient detail for these applications (for example, to identify the full automorphism groups of these combinatorial structures) requires thorough knowledge of the set of primitive and quasiprimitive overgroups of a given quasiprimitive subgroup of the full symmetric group, a problem addressed in work of the first author and R. W. Baddeley (Baddeley & Praeger 2003). If such a subgroup preserves a non-trivial cartesian decomposition, then the full stabiliser of this decomposition, which is a wreath product in product action, is a natural overgroup. It turns out that the best way to find such overgroups is to locate the corresponding invariant cartesian decompositions.

Quasiprimitivity may be equivalently defined for finite groups as the requirement that all minimal normal subgroups are transitive. Weakening this to the requirement that at least one minimal normal subgroup is transitive gives the larger family of innately transitive groups. An 'O'Nan–Scott theory' of finite innately transitive permutation groups was developed by Bamberg and the first author (Bamberg 2003, Bamberg & Praeger 2004). Applications of this theory also require detailed knowledge of homogeneous cartesian decompositions left invariant by these groups.

Our current knowledge of finite primitive permutation groups is very strong, because not only do we have the O'Nan–Scott Theorem, which divides the primitive permutation groups into families according to the action and the structure of the socle, but we also have available a description which divides the primitive permutation groups into families according to the nature of their primitive overgroups, and most significantly we know that these two descriptions are essentially the same (Praeger 1990). Information concerning both the socle and all overgroups is central to our understanding of primitive permutation groups. This book attempts to advance our understanding of innately transitive and quasiprimitive permutation groups in a similar fashion by analysing such groups from both viewpoints and we do this for the family of (possibly infinite) groups with an FCR-plinth. Indeed, the quasiprimitive version of the O'Nan–Scott Theorem (Praeger 1993) reaches its weakest conclusions when considering finite quasiprimitive permutation groups G whose socle M is a minimal normal subgroup of G and is such that a point stabiliser M_α in M is neither trivial nor projects onto each minimal normal subgroup of M (see Section 7.5). (This is case III(b)(i) in the terminology of (Praeger 1993), and case PA in (Baddeley & Praeger 2003).) For primitive groups the analogous case corresponds essentially to those primitive permutation groups that arise as subgroups of wreath products in product action. It is to be hoped that the overgroup viewpoint can shed more light on this problematic case.

The theory of invariant cartesian decompositions is also useful for studying certain problems in geometry and combinatorics. For example, Baumeister (Baumeister 1997) determined the class of finite two-dimensional grids in which the stabiliser of one of the parallel classes of lines acts primitively on this class. Applying the theory of cartesian decompositions can give a new proof of Baumeister's results, and can extend them to results on flag-transitive multi-dimensional grids with an automorphism group innately transitive on the vertices.

Wreath products and cartesian decompositions also play a role in the reduction theorem for finite primitive distance transitive graphs in (Praeger, Saxl & Yokoyama 1987). The automorphism group of such a graph is either almost simple, or affine, or a wreath product in product action preserving a cartesian decomposition (as, for instance, with the