The geometry of diagonal groups

Peter J. Cameron
University of St Andrews

UWA Groups and Combinatorics Seminar
22 May 2020

Background 1

Cheryl Praeger and Csaba Schneider started this research some time ago.

In Shenzhen in 2018, they invited Rosemary Bailey and me to join them.

Background 2

Things went on slowly, but at the six-month programme on groups at the Isaac Newton Institute in Cambridge, we hoped to bring it to a conclusion.

But the coronavirus had other ideas. So we put it on hold and all went home.

Background 3

Rosemary and I were in the fortunate position that we could continue working on the problem in the old-fashioned way. I had a vision of what the main result should be, and, thanks very largely to Rosemary, we were able to carry this through. So, in this talk, I want to tell you about how to recognise structures which admit diagonal groups as groups of automorphisms.

The motivating example comes from the 2-dimensional case. If T is a group and $D(T,2)$ the 2-dimensional diagonal group built from T, then $D(T,2)$ is precisely the automorphism group of the Latin square graph associated with the Cayley table of T. This is joint work of the four of us (RAB, PJC, CEP and CS).

The case $m = 1$

I will not be talking about this case, but I mention it now.
In 1968, Peter Neumann, Charles Sims and James Wiegold published a paper with the wonderful title "Counterexamples to a theorem of Cauchy".
Cauchy had "proved" that, if a primitive group has degree a prime number plus one, then the group must be doubly transitive. Neumann, Sims and Wiegold pointed out that, if T is a finite simple group and $G = T \times T$ acts on T, the first factor by left multiplication by the inverse, the second by right multiplication, then G is primitive but not doubly transitive.
There are many simple groups whose order is one more than a prime, starting with A_5.
A challenge to number theorists: Are there infinitely many finite simple groups which give counterexamples to Cauchy's theorem in this way?

The O'Nan–Scott Theorem

Theorem

A finite primitive permutation group is of one of the following types: affine, wreath product, diagonal, or almost simple.

Affine groups preserve affine spaces; wreath products preserve Cartesian structures (as we discuss later); almost simple groups form a ragbag, and there is no hope for a uniform description of the structures they act on.

Our aim is to understand the geometric structure underlying diagonal groups. But, unlike in the O'Nan–Scott theorem, we do not assume that these groups are finite or primitive.

Diagonal groups, 1

Let T be a group, and m a positive integer. The diagonal group $D(T,m)$ has a normal subgroup of the following shape: take the group T^{m+1} (factors numbered from 0 to m) acting on the set Ω of right cosets of its diagonal subgroup $\{(t,t,\ldots,t) : t \in T\}$.
This group is normalised by further transformations: $\mathrm{Aut}(T)$, acting in the same way on all coordinates; and the symmetric group S_{m+1}, permuting the coordinates.
It is more convenient (though we lose some symmetry) to take a different representation. Every coset of the diagonal group has a unique representative with the identity in coordinate 0. We denote this representative by $[t_1,\ldots,t_m]$.

Diagonal groups, 2

- The factors $T_1, \ldots, T_m$ act as usual by right multiplication:

$$x \in T_1 : [t_1, t_2 \ldots, t_m] \mapsto [t_1 x, t_2, \ldots, t_m].$$

- T_0 acts by simultaneous left multiplication of all coordinates by the inverse:

$$x \in T_0 : [t_1, t_2, \ldots, t_m] \mapsto [x^{-1} t_1, x^{-1} t_2, \ldots, x^{-1} t_m].$$

- Automorphisms act in the same way on all coordinates.
- S_m acts by permuting the coordinates.
- The transposition τ of coordinates 0 and 1 acts as

$$\tau : [t_1, t_2, \ldots, t_m] \mapsto [t_1^{-1}, t_1^{-1} t_2, \ldots, t_1^{-1} t_m].$$

Diagonal groups, 3

This defines the group $D(T,m)$ as a permutation group on T^m for any group T and positive integer m. The "simple diagonal" type in the O'Nan–Scott Theorem is obtained when T is a finite simple group.

As said before, I will not be discussing the case $m = 1$, where there is no geometry to help us; we assume that $m \geq 2$.

Our goal in setting out was the following:

> *Give a combinatorial description of a structure with automorphism group induced by $D(T,m)$, so that the group T does not have to be built in to the construction, but emerges naturally from the combinatorics.*

In projective geometry, the projective planes are "wild", but higher dimensional spaces are coordinatised by a division ring. We will see a similar phenomenon here.

Partitions

A partition of Ω can be thought of in any of three ways:

- a set of non-empty, pairwise disjoint subsets of Ω whose union is Ω;
- the set of equivalence classes of an equivalence relation on Ω;
- the kernel of a function F on Ω, that is, the set of inverse images of points in the range of F.

The set $\mathbb{P}(\Omega)$ of partitions of Ω is partially ordered by refinement: $P \preceq Q$ if every part of P is contained in a part of Q.

The partition lattice

With this order, $\mathbb{P}(\Omega)$ is a lattice: any two partitions P and Q have a unique infimum or meet $P \wedge Q$, and a unique supremum or join $P \vee Q$.

- $P \wedge Q$ is the partition of Ω whose parts are all *non-empty* intersections of a part of P and a part of Q.
- $P \vee Q$ is the partition into connected components of the graph in which two points are adjacent if they lie in the same part of *either* P or Q.

A subset of $\mathbb{P}(\Omega)$ is a sublattice if it is closed under the meet and join operations of $\mathbb{P}(\Omega)$.
We also require the notion of a join-semilattice, closed under join but maybe not under meet.

Coset partitions

Let G be a finite group. For each subgroup H of G, consider the partition P_H of G into right cosets of H. We call this a coset partition.

Now, if H and K are subgroups of G, then we have

- $P_H \preceq P_K$ if and only if $H \leq K$;
- $P_H \wedge P_K = P_{H \cap K}$;
- $P_H \vee P_K = P_{\langle H,K \rangle}$.

So the collection of all coset partitions of G forms a sublattice of $\mathbb{P}(G)$ which is isomorphic to the subgroup lattice of G, under the map $H \mapsto P_H$.

Cartesian decompositions

If you want to know everything about these, here is the place to look:

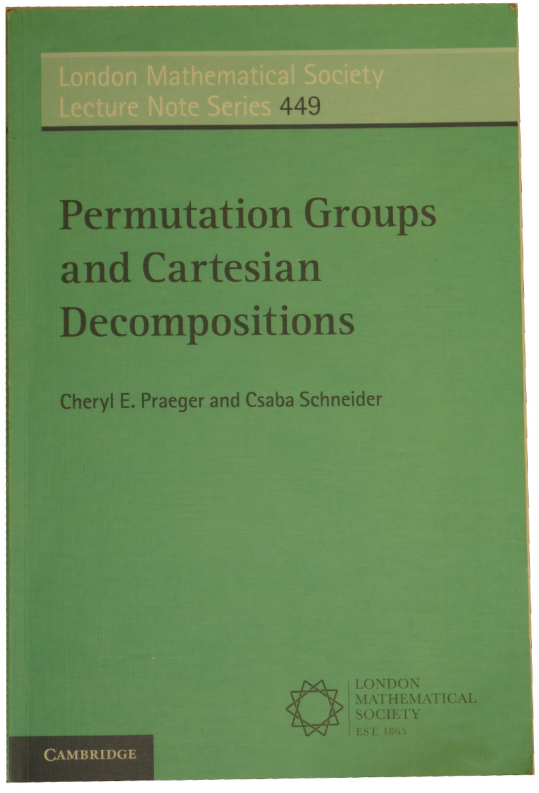

I will tell you just what you need here.

Cartesian lattices, 1

Suppose that $\Omega = A^n$ for some alphabet A, with $|A| > 1$. For $I \subseteq \{1, \ldots, n\}$, let Q_I be the partition of Ω defined by the equivalence relation $\equiv_I$, where $a \equiv_I b$ if $a_j = b_j$ for all $j \notin I$.
The partitions Q_I for $I \subseteq \{1, \ldots, n\}$ form a sublattice of $\mathbb{P}(\Omega)$; indeed, the map $I \mapsto Q_I$ is an isomorphism from the Boolean lattice of subsets of $\{1, \ldots, n\}$ to this sublattice. We call such a lattice a Cartesian lattice on Ω. Its dimension is defined to be n.
An alternative approach uses the Hamming graph $\mathrm{Ham}(n, A)$, with vertex set A^n, two vertices joined if they have Hamming distance 1, that is, they differ in just one coordinate.
The minimal (non-trivial) elements of the Cartesian lattice have the form Q_i, which is a partition into maximal cliques of the Hamming graph. So the Cartesian lattice and the Hamming graph determine each other.

Cartesian lattices, 2

In view of the preceding remark, the Hamming graph and the Cartesian lattice have the same automorphism group, namely, the wreath product $\mathrm{Sym}(A)$ Wr S_n.
We take Cartesian lattices to be the geometries associated with wreath products in the O'Nan–Scott theorem.
It is perhaps worth noting that the notions of Cartesian lattice and Hamming graph can be generalised, to the mixed alphabet case: we take $\Omega = A_1 \times \cdots \times A_n$, and proceed in the same way. (In this case the automorphism group is not the wreath product as above, since coordinates associated with alphabets of different sizes are not interchangeable.)
However, this complication does not arise in what we are going to do.

Latin squares, 1

You probably think of a Latin square as something like this: a square array of size $n \times n$ filled with letters from an alphabet of size n, so that each letter occurs once in each row and column.

A	B	C
B	C	A
C	A	B

We are going to give a different definition. Let Ω consist of the n^2 cells of the array. We have three partitions of Ω: R, the rows; C, the columns; and L, the letters (the partition into sets of cells containing the same letter).

Latin squares, 2

A	B	C
B	C	A
C	A	B

1	2	3
4	5	6
7	8	9

- $R = \{\{1,2,3\},\{4,5,6\},\{7,8,9\}\}$;
- $C = \{\{1,4,7\},\{2,5,8\},\{3,6,9\}\}$;
- $L = \{\{1,6,8\},\{2,4,9\},\{3,5,7\}\}$.

Together with E (the partition into singletons) and U (the partition with a single part), these three partitions form a lattice. It has the very special property that, if one of R, C, L is omitted, the resulting four partitions form a Cartesian lattice on Ω.

This property characterises Latin squares.

Latin squares, 3

With the partition definition, we could define an automorphism of a Latin square to be a permutation of Ω fixing $\{R, C, L\}$ setwise. (These mappings are usually called paratopisms in the Latin squares literature.)
Latin squares are extremely prolific. It is known that the number of Latin squares of order n grows faster than exponentially in n^2. Moreover, almost all of them have only trivial paratopism group.
However, one case is interesting to us: the Cayley table of a group T is a Latin square, and its paratopism group is the diagonal group $D(T, 2)$ defined earlier.

Diagonal semilattices

Let us return to diagonal groups for a moment. Recall that $D(T,m)$ acts on T^m, where m factors of the base group T^{m+1} each act on the corresponding coordinate of T^m by right multiplication, while the last factor acts by simultaneous left multiplication by the inverse.
Let $T_0,\dots,T_m$ be these subgroups, and let Q_i be the partition of T^m into cosets of T_i for $i=0,1,\dots,m$.
The join-semilattice generated by $Q_0,\dots,Q_m$ (it is not a lattice for $m\geq 3$) is an object which we will call a diagonal semilattice and denote by $\mathcal{D}(T,m)$.
Its automorphism group is the diagonal group $D(T,m)$.

The main theorem

Theorem

Let $m \geq 2$, and let $Q_0, Q_1, \ldots, Q_m$ be partitions of Ω. Suppose that any m of these partitions generate an m-dimensional Cartesian lattice on Ω, in which the given partitions are the minimal (non-trivial) elements.

- *If $m = 2$, then $\{Q_0, Q_1, Q_2\}$, together with E and U, form a Latin square, unique up to isotopism; every Latin square arises in this way.*
- *If $m \geq 3$, then there is a group T, determined up to isomorphism, such that the join-semilattice generated by $\{Q_0, \ldots, Q_m\}$ is the diagonal semilattice $\mathcal{D}(T, m)$.*

As promised, for $m = 2$ the situation is chaotic, but for $m \geq 3$ the algebraic structure coordinatising the semilattice (the group T) emerges naturally from the combinatorics.

The proof

Time does not permit me to go through the whole proof. I hope that this sketch will suffice.

First, there is nothing to prove if $m = 2$. Two of the partitions give Ω the structure of a square grid, and the third partitions it into the positions of letters; the remaining conditions force the letters to form a Latin square.

Also, the proof for $m \geq 3$ is by induction on m. I will not give any details about the induction step; the crucial part is starting the induction at $m = 3$, where groups first make their appearance.

Latin cubes, 1

Our hypotheses for $m = 3$ are equivalent to a certain kind of Latin cube.
Unfortunately there are several incompatible definitions of Latin cubes in the literature. I will stick to the one we need.

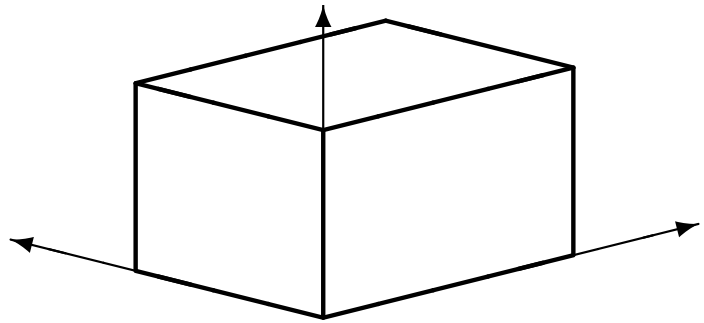

Think of a combinatorial cube (3-dimensional array). We shall regard it as being given by a Cartesian lattice. We call a slice parallel to one of the coordinate planes a layer, and a line parallel to one of the coordinate axes a line.
Thus, if Q_1, Q_2, Q_3 are the minimal partitions generating the Cartesian lattice, then lines are parts of one of these three partitions, while layers are parts of one of the partitions $Q_{12} = Q_1 \vee Q_2$, etc.

Latin cubes, 2

Now a Latin cube (of the sort we want) has an additional partition Q_0, which we think of as corresponding to letters written in the cells of the array. We require that

- each letter occurs exactly once in each layer.

We require another property, which doesn't have a standard name, which we call regularity:

- in any two parallel lines, the sets of letters occurring in cells in those lines are either equal or disjoint.

Most of our proof involves giving a description of all regular Latin cubes in the above sense: these turn out to be equivalent to the case $m = 3$ of the main theorem.

Latin squares and quasigroups

Given a Latin square, we can label the rows, columns and letters by elements of a set T; then the given square is the Cayley table of a quasigroup on the set T.
Changing the labellings corresponds to an isotopism of the quasigroup.
A Latin square is isomorphic to the Cayley table of a group if and only if it satisfies the quadrangle condition, which I now describe. Given two rows and two columns, there are four letters lying in the positions of these rows and columns. Suppose that another choice of two rows and two columns gives rise to another set of four letters. We say that the quadrangle condition holds if, whenever three of the four letters coincide, then the fourth does also, for all possible choices. The following result goes back to Frolov (1890):

Theorem

A quasigroup is isotopic to a group if and only if its Cayley table satisfies the quadrangle condition.

Quasigroups from Latin cubes

Let $Q_0, \ldots, Q_3$ be partitions of Ω satisfying the hypotheses of the main theorem. Any three of them generate a Cartesian lattice, and the fourth gives it the structure of a regular Latin cube.

We obtain Latin squares as quotients. Pick one of the partitions Q_i. Let $\overline{\Omega}$ be the set of parts of Q_i, and for $j \neq i$, let $\overline{Q_j}$ be the partition of $\overline{\Omega}$ corresponding to the partition $Q_i \vee Q_j$ of Ω.

These partitions form a Latin square.

Thus we have four Latin squares, and each allows three choices of labelling; but because of interdependencies, there are six labellings which can be chosen independently.

The heart of the proof

We have to show three things:

- these Latin squares satisfy the quadrangle condition, and hence are isotopic to groups;
- we have six labellings to choose, and this can be done in such a way that the four groups are all isomorphic, and the Latin squares are Cayley tables with (x, y) entry $x^{-1}y$ (this is not the usual Cayley table, but is isotopic to it);
- using this, Ω can be identified with the diagonal semilattice $\mathcal{D}(T, 3)$ over the group T.

At this point, the proof for $m = 3$ is complete.

The diagonal graph

The diagonal graph stands in a similar relation to a diagonal semilattice as the Hamming graph does to a Cartesian lattice. Let Ω be a set, and $Q_0, \ldots, Q_m$ partitions of Ω satisfying the hypotheses of our main theorem: thus any m of them generate a Cartesian lattice.

The diagonal graph corresponding to these data has vertex set Ω, with two vertices adjacent if and only if they lie in the same part of Q_i for some (unique) value of i.

Thus the edges of the graph fall into $m+1$ types, so that any m of these types comprise a Hamming graph $\mathrm{Ham}(m, |T|)$.

For $m > 2$, the diagonal semilattice can be recovered from the graph, since its maximal cliques are the parts of the partitions Q_i, and we can decide combinatorially whether two maximal cliques are of the same type. So the automorphism group of the graph is the diagonal group $D(T, m)$. We denote the graph by $\Gamma(T, m)$.

Some examples

In the case $m = 2$, the graph is precisely the Latin square graph associated with the Latin square. We ignore this case and assume that $m \geq 3$.

Consider the case $|T| = 2$. In this case, the Hamming graph is the m-dimensional cube. The extra edges we have to add to get the diagonal graph are precisely those joining antipodal vertices of the cube. So the graph in this case is the folded cube, and is distance-transitive.

This provides a good model for thinking about these graphs: the diagonal graph consists of the Hamming graph with some extra edges joining vertices at maximal distance.

The Hall–Paige conjecture

A complete mapping of a group G is a bijection $\phi : G \to G$ such that the map $\psi : G \to G$ defined by $\psi(x) = x\phi(x)$ is also a bijection. It is well known that the following three conditions on a group G are equivalent:

- G has a complete mapping;
- the Cayley table of G has a transversal;
- the Cayley table of G has an orthogonal mate.

In 1955, Hall and Paige showed that a group with these properties has trivial or non-cyclic Sylow 2-subgroups, and conjectured the converse. This conjecture was proved by Wilcox, Evans and Bray in 2009, using the Classification of Finite Simple Groups.

Clique number and chromatic number

In connection with synchronization, Bray *et al.* showed that, if T is a finite simple group, and $m \geq 2$, then the clique number and chromatic number of the diagonal graph are both equal to $|T|$. We have a conjectured extension to arbitrary finite groups, not entirely proved yet.

Theorem

Let Γ be the graph $\Gamma(T, m)$, where $m \geq 2$ and T is a finite group.

- *The clique number of Γ is $|T|$.*
- *If m is odd, then the chromatic number of Γ is $|T|$.*
- *If m is even and T has a complete mapping, then the chromatic number of Γ is $|T|$.*

Problem

Is it true that, for m even, the chromatic number of Γ is equal to $|T|$ if and only if T has a complete mapping?

References

- J. N. Bray, Q. Cai, P. J. Cameron, P. Spiga and H. Zhang, The Hall–Paige conjecture, and synchronization for affine and diagonal groups, *J. Algebra* **545** (2020), 27–42; doi: 10.1016/j.jalgebra.2019.02.025
- C. E. Praeger and C. Schneider, *Permutation Groups and Cartesian Decompositions*, London Math. Soc. Lecture Notes **449**, Cambridge Univ. Press 2018; doi: 10.1017/9781139194006