The geometry of diagonal groups

Peter J. Cameron, University of St Andrews

53rd Southeastern Conference on Combinatorics,
Graph Theory and Computing
9 March 2022

Joint work with Rosemary Bailey, Michael Kinyon,
Cheryl Praeger and Csaba Schneider

The theorem

In the first part of the talk, I will describe our theorem. In the second part, time permitting, I will talk about some extensions and applications.

An analogy

I begin with an analogy. If you know any projective geometry, you will be aware of the following phenomenon:

An analogy

I begin with an analogy. If you know any projective geometry, you will be aware of the following phenomenon:

- a 1-dimensional projective geometry (a projective line) has no incidence structure at all; it is just a set.

An analogy

I begin with an analogy. If you know any projective geometry, you will be aware of the following phenomenon:

- a 1-dimensional projective geometry (a projective line) has no incidence structure at all; it is just a set.
- 2-dimensional projective geometries (projective planes) exist in wild profusion, so that there is no hope of classification.

An analogy

I begin with an analogy. If you know any projective geometry, you will be aware of the following phenomenon:

- a 1-dimensional projective geometry (a projective line) has no incidence structure at all; it is just a set.
- 2-dimensional projective geometries (projective planes) exist in wild profusion, so that there is no hope of classification.
- For higher dimensions, a projective geometry is highly structured, and is coordinatised by an algebraic object (a division ring).

An analogy

I begin with an analogy. If you know any projective geometry, you will be aware of the following phenomenon:

- a 1-dimensional projective geometry (a projective line) has no incidence structure at all; it is just a set.
- 2-dimensional projective geometries (projective planes) exist in wild profusion, so that there is no hope of classification.
- For higher dimensions, a projective geometry is highly structured, and is coordinatised by an algebraic object (a division ring).

I am going to show you a very similar phenomenon: "wild profusion" will mean arbitrary Latin squares, while the "algebraic object" will be a group.

Permutation groups

I begin with a little diversion into permutation group theory. G will denote a permutation group on Ω.

Permutation groups

I begin with a little diversion into permutation group theory. G will denote a permutation group on Ω.
G is transitive if no non-trivial subset of Ω is G-invariant; it is primitive if no non-trivial partition of Ω is G-invariant. (The trivial subsets or partitions are those invariant under the symmetric group on Ω.

Permutation groups

I begin with a little diversion into permutation group theory. G will denote a permutation group on Ω.

G is transitive if no non-trivial subset of Ω is G-invariant; it is primitive if no non-trivial partition of Ω is G-invariant. (The trivial subsets or partitions are those invariant under the symmetric group on Ω.

Many (but not all) questions about permutation groups can be reduced to the case where the group is primitive. This has been a standard technique since Jordan in the 19th century.

Permutation groups

I begin with a little diversion into permutation group theory. G will denote a permutation group on Ω.
G is transitive if no non-trivial subset of Ω is G-invariant; it is primitive if no non-trivial partition of Ω is G-invariant. (The trivial subsets or partitions are those invariant under the symmetric group on Ω.
Many (but not all) questions about permutation groups can be reduced to the case where the group is primitive. This has been a standard technique since Jordan in the 19th century.
So how can we understand primitive groups?

The O'Nan–Scott Theorem

The following theorem (and indeed rather more) was proved by Michael O'Nan and Leonard Scott (independently) in 1979. The version I need here is little more than is in Jordan's *Traité des Substitutions*.

The O'Nan–Scott Theorem

The following theorem (and indeed rather more) was proved by Michael O'Nan and Leonard Scott (independently) in 1979. The version I need here is little more than is in Jordan's *Traité des Substitutions*.

Theorem

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

The O'Nan–Scott Theorem

The following theorem (and indeed rather more) was proved by Michael O'Nan and Leonard Scott (independently) in 1979. The version I need here is little more than is in Jordan's *Traité des Substitutions*.

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 I discuss later); almost simple groups form a ragbag, and there is no hope for a uniform description of the structures they act on.

The O'Nan–Scott Theorem

The following theorem (and indeed rather more) was proved by Michael O'Nan and Leonard Scott (independently) in 1979. The version I need here is little more than is in Jordan's *Traité des Substitutions*.

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 I 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.

Background 1

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

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 in 2020, we hoped to bring it to a conclusion.

Background 2

Things went on slowly, but at the six-month programme on groups at the Isaac Newton Institute in Cambridge in 2020, we hoped to bring it to a conclusion.

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

Diagonal groups

Let m be a positive integer and T a group, finite or infinite. I define the diagonal group $D(T,m)$ to be the group of permutations of $\Omega = T^m$ generated by the following transformations. (I put the elements of Ω in square brackets to distinguish them from group elements.)

Diagonal groups

Let m be a positive integer and T a group, finite or infinite. I define the diagonal group $D(T,m)$ to be the group of permutations of $\Omega = T^m$ generated by the following transformations. (I put the elements of Ω in square brackets to distinguish them from group elements.)

- The group T^m acting by right multiplication.

Diagonal groups

Let m be a positive integer and T a group, finite or infinite. I define the diagonal group $D(T, m)$ to be the group of permutations of $\Omega = T^m$ generated by the following transformations. (I put the elements of Ω in square brackets to distinguish them from group elements.)

- The group T^m acting by right multiplication.
- another copy T_0 of T acting by simultaneous left multiplication of all coordinates by the inverse.

Diagonal groups

Let m be a positive integer and T a group, finite or infinite. I define the diagonal group $D(T,m)$ to be the group of permutations of $\Omega = T^m$ generated by the following transformations. (I put the elements of Ω in square brackets to distinguish them from group elements.)

- The group T^m acting by right multiplication.
- another copy T_0 of T acting by simultaneous left multiplication of all coordinates by the inverse.
- $\mathrm{Aut}(T)$ acting in the same way on all coordinates.

Diagonal groups

Let m be a positive integer and T a group, finite or infinite. I define the diagonal group $D(T, m)$ to be the group of permutations of $\Omega = T^m$ generated by the following transformations. (I put the elements of Ω in square brackets to distinguish them from group elements.)

- The group T^m acting by right multiplication.
- another copy T_0 of T acting by simultaneous left multiplication of all coordinates by the inverse.
- $\mathrm{Aut}(T)$ acting in the same way on all coordinates.
- S_m acting by permuting the coordinates.

Diagonal groups

Let m be a positive integer and T a group, finite or infinite. I define the diagonal group $D(T,m)$ to be the group of permutations of $\Omega = T^m$ generated by the following transformations. (I put the elements of Ω in square brackets to distinguish them from group elements.)

- The group T^m acting by right multiplication.
- another copy T_0 of T acting by simultaneous left multiplication of all coordinates by the inverse.
- $\mathrm{Aut}(T)$ acting in the same way on all coordinates.
- S_m acting by permuting the coordinates.
- An element τ:

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

Diagonal groups

Let m be a positive integer and T a group, finite or infinite. I define the diagonal group $D(T, m)$ to be the group of permutations of $\Omega = T^m$ generated by the following transformations. (I put the elements of Ω in square brackets to distinguish them from group elements.)

- The group T^m acting by right multiplication.
- another copy T_0 of T acting by simultaneous left multiplication of all coordinates by the inverse.
- $\mathrm{Aut}(T)$ acting in the same way on all coordinates.
- S_m acting by permuting the coordinates.
- An element τ:

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

Don't remember the details: this is just a group built from T and m.

Partitions

Our geometry will be defined in terms of partitions. So here is a brief introduction.

Partitions

Our geometry will be defined in terms of partitions. So here is a brief introduction.

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;

Partitions

Our geometry will be defined in terms of partitions. So here is a brief introduction.

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;

Partitions

Our geometry will be defined in terms of partitions. So here is a brief introduction.

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.

Partitions

Our geometry will be defined in terms of partitions. So here is a brief introduction.

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: $P \preccurlyeq Q$ 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 $P \vee Q$.

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.

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.

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 $\mathbb{P}(\Omega)$.

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.

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

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

- $P_H \preccurlyeq P_K$ if and only if $H \leqslant K$;

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

- $P_H \preccurlyeq P_K$ if and only if $H \leqslant K$;
- $P_H \wedge P_K = P_{H \cap K}$;

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

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

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

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

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 $H \mapsto P_H$.

Structures for wreath products

These have many different descriptions. Praeger and Schneider, who discussed them before moving on to diagonal groups, call them Cartesian decompositions.

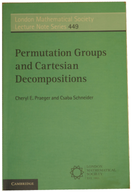

Structures for wreath products

These have many different descriptions. Praeger and Schneider, who discussed them before moving on to diagonal groups, call them Cartesian decompositions.

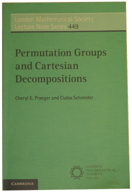

Graph theorists call them Hamming graphs. The name hints at a connection with coding theory. Indeed, Delsarte called them Hamming schemes. This description, however, loses the order relation. Statisticians call them completely crossed orthogonal block structures.

Structures for wreath products

These have many different descriptions. Praeger and Schneider, who discussed them before moving on to diagonal groups, call them Cartesian decompositions.

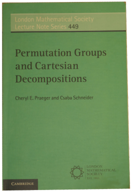

Graph theorists call them Hamming graphs. The name hints at a connection with coding theory. Indeed, Delsarte called them Hamming schemes. This description, however, loses the order relation. Statisticians call them completely crossed orthogonal block structures.

I will use the term Cartesian lattices.

Cartesian lattices

The Boolean lattice $\mathcal{B}_n$ is the lattice of all subsets of $\{1, \ldots, n\}$.

Cartesian lattices

The Boolean lattice $\mathcal{B}_n$ is the lattice of all subsets of $\{1, \ldots, n\}$. Let A be an alphabet, finite or infinite (with $|A| > 1$). Let $\Omega = A^n$ be the set of all words of length n over the alphabet A.

Cartesian lattices

The Boolean lattice $\mathcal{B}_n$ is the lattice of all subsets of $\{1,\ldots,n\}$. Let A be an alphabet, finite or infinite (with $|A| > 1$). Let $\Omega = A^n$ be the set of all words of length n over the alphabet A. For $I \subseteq \{1,\ldots,n\}$, let Q_I be the partition of Ω corresponding to the equivalence relation $\equiv_I$, where

$$(a_1,\ldots,a_n) \equiv_I (b_1,\ldots,b_n) \Leftrightarrow (\forall j \notin I)(a_j = b_j).$$

Cartesian lattices

The Boolean lattice $\mathcal{B}_n$ is the lattice of all subsets of $\{1,\ldots,n\}$. Let A be an alphabet, finite or infinite (with $|A| > 1$). Let $\Omega = A^n$ be the set of all words of length n over the alphabet A. For $I \subseteq \{1,\ldots,n\}$, let Q_I be the partition of Ω corresponding to the equivalence relation $\equiv_I$, where

$$(a_1,\ldots,a_n) \equiv_I (b_1,\ldots,b_n) \Leftrightarrow (\forall j \notin I)(a_j = b_j).$$

Now the partitions Q_I for $I \subseteq \{1,\ldots,n\}$ form a sublattice of the partition lattice on Ω which is isomorphic to $\mathcal{B}_n$ by the map $I \mapsto Q_I$.

Cartesian lattices

The Boolean lattice $\mathcal{B}_n$ is the lattice of all subsets of $\{1, \ldots, n\}$. Let A be an alphabet, finite or infinite (with $|A| > 1$). Let $\Omega = A^n$ be the set of all words of length n over the alphabet A. For $I \subseteq \{1, \ldots, n\}$, let Q_I be the partition of Ω corresponding to the equivalence relation $\equiv_I$, where

$$(a_1, \ldots, a_n) \equiv_I (b_1, \ldots, b_n) \Leftrightarrow (\forall j \notin I)(a_j = b_j).$$

Now the partitions Q_I for $I \subseteq \{1, \ldots, n\}$ form a sublattice of the partition lattice on Ω which is isomorphic to $\mathcal{B}_n$ by the map $I \mapsto Q_I$.
I will call this a Cartesian lattice. Note that the group of permutations of Ω mapping the lattice to itself (as set of partitions) is the wreath product $\mathrm{Sym}(A) \operatorname{Wr} \mathrm{Sym}(\{1, \ldots, n\})$.

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.

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

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

Latin squares exist in great profusion. There are more than $\exp(m^2)$ Latin squares of order m; exact numbers are only known up to $m = 11$.

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

Latin squares exist in great profusion. There are more than $\exp(m^2)$ Latin squares of order m; exact numbers are only known up to $m = 11$.
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

Latin squares, 2

A	B	C
B	C	A
C	A	B

1	2	3
4	5	6
7	8	9

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

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 2-dimensional Cartesian lattice on Ω.

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 2-dimensional 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, 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.)

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. (This fact is maybe not as well known as it should be!)

Diagonal semilattices

Let us return to diagonal groups for a moment. Recall that $D(T,m)$ acts on T^m, where m copies $T_1,\ldots,T_m$ of T act on the corresponding coordinate of T^m by right multiplication, while the last factor T_0 acts by simultaneous left multiplication by the inverse.

Diagonal semilattices

Let us return to diagonal groups for a moment. Recall that $D(T,m)$ acts on T^m, where m copies $T_1,\ldots,T_m$ of T act on the corresponding coordinate of T^m by right multiplication, while the last factor T_0 acts by simultaneous left multiplication by the inverse.

Let $Q_0,\ldots,Q_m$ be the orbit partitions of $\Omega = T^m$ corresponding to these groups. Thinking of T^m as a group, these are the coordinate partitions of the coordinate groups $T_1,\ldots,T_m$ and the diagonal subgroup of T^m (hence the name).

Diagonal semilattices

Let us return to diagonal groups for a moment. Recall that $D(T,m)$ acts on T^m, where m copies $T_1,\ldots,T_m$ of T act on the corresponding coordinate of T^m by right multiplication, while the last factor T_0 acts by simultaneous left multiplication by the inverse.

Let $Q_0,\ldots,Q_m$ be the orbit partitions of $\Omega = T^m$ corresponding to these groups. Thinking of T^m as a group, these are the coordinate partitions of the coordinate groups $T_1,\ldots,T_m$ and the diagonal subgroup of T^m (hence the name).

The join-semilattice generated by $Q_0,\ldots,Q_m$ (it is not a lattice for $m \geqslant 3$) is an object which we will call a diagonal semilattice and denote by $\mathcal{D}(T,m)$.

Diagonal semilattices

Let us return to diagonal groups for a moment. Recall that $D(T,m)$ acts on T^m, where m copies $T_1,\ldots,T_m$ of T act on the corresponding coordinate of T^m by right multiplication, while the last factor T_0 acts by simultaneous left multiplication by the inverse.

Let $Q_0,\ldots,Q_m$ be the orbit partitions of $\Omega = T^m$ corresponding to these groups. Thinking of T^m as a group, these are the coordinate partitions of the coordinate groups $T_1,\ldots,T_m$ and the diagonal subgroup of T^m (hence the name).

The join-semilattice generated by $Q_0,\ldots,Q_m$ (it is not a lattice for $m \geqslant 3$) is an object which we will call a diagonal semilattice and denote by $\mathcal{D}(T,m)$.

Theorem

The automorphism group of $\mathcal{D}(T,m)$ is the diagonal group $D(T,m)$.

The main theorem

Theorem

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

The main theorem

Theorem

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

- *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.*

The main theorem

Theorem

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

- *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 $m \geqslant 3$, 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 $\mathcal{D}(T, m)$.*

The main theorem

Theorem

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

- *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 $m \geqslant 3$, 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 $m \geqslant 3$ the algebraic structure coordinatising the semilattice (the group T) emerges naturally from the combinatorics.

The main theorem

Theorem

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

- *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 $m \geqslant 3$, 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 $m \geqslant 3$ the algebraic structure coordinatising the semilattice (the group T) emerges naturally from the combinatorics.
I am very proud of the proof we found, but I fear I don't have time even for a sketch.

Applications

In the remaining time I will briefly mention some applications and extensions of this result.

The diagonal graph

There is a close connection between the Cartesian lattice and the Hamming graph. Recall that A^n is the set of words of length n over the alphabet A. The Hamming graph has vertex set A^n; two vertices are joined if as words the agree in all positions except one (that is, they have Hamming distance 1).

The diagonal graph

There is a close connection between the Cartesian lattice and the Hamming graph. Recall that A^n is the set of words of length n over the alphabet A. The Hamming graph has vertex set A^n; two vertices are joined if as words the agree in all positions except one (that is, they have Hamming distance 1).
Said otherwise, two elements of A^n are joined if they are contained in the same part of a minimal non-trivial partition of the Cartesian lattice. So graph and lattice determine each other.

The diagonal graph

There is a close connection between the Cartesian lattice and the Hamming graph. Recall that A^n is the set of words of length n over the alphabet A. The Hamming graph has vertex set A^n; two vertices are joined if as words the agree in all positions except one (that is, they have Hamming distance 1).
Said otherwise, two elements of A^n are joined if they are contained in the same part of a minimal non-trivial partition of the Cartesian lattice. So graph and lattice determine each other.
In a similar way, we can construct a graph from the diagonal semilattice: two vertices are joined if they are contained in the same part of a minimal non-trivial partition of $\mathcal{D}(T,m)$.

This is an interesting graph, worth further investigation. I mention a few things about it.

This is an interesting graph, worth further investigation. I mention a few things about it.

- Except for a few very small cases, the semilattice can be reconstructed from the graph; so its automorphism group is the diagonal group $D(T, m)$.

This is an interesting graph, worth further investigation. I mention a few things about it.

- Except for a few very small cases, the semilattice can be reconstructed from the graph; so its automorphism group is the diagonal group $D(T, m)$.
- For $m = 2$, it is a (strongly regular) Latin square graph, while for $|T| = 2$, it is a (distance-transitive) folded cube.

This is an interesting graph, worth further investigation. I mention a few things about it.

- Except for a few very small cases, the semilattice can be reconstructed from the graph; so its automorphism group is the diagonal group $D(T, m)$.
- For $m = 2$, it is a (strongly regular) Latin square graph, while for $|T| = 2$, it is a (distance-transitive) folded cube.
- Except for a few very small cases, its clique number is $|T|$.

This is an interesting graph, worth further investigation. I mention a few things about it.

- Except for a few very small cases, the semilattice can be reconstructed from the graph; so its automorphism group is the diagonal group $D(T, m)$.
- For $m = 2$, it is a (strongly regular) Latin square graph, while for $|T| = 2$, it is a (distance-transitive) folded cube.
- Except for a few very small cases, its clique number is $|T|$.
- If m is odd, or if $|T|$ is odd, or if the Sylow 2-subgroups of T are non-cyclic, its chromatic number is also $|T|$.

The chromatic number mentioned in the fourth point above depends on the truth of the Hall–Paige conjecture on complete mappings of groups, whose proof depends on the classification of finite simple groups.

The chromatic number mentioned in the fourth point above depends on the truth of the Hall–Paige conjecture on complete mappings of groups, whose proof depends on the classification of finite simple groups.
This has an application to the question of synchronization of finite automata, about which I spoke here in 2013.

The chromatic number mentioned in the fourth point above depends on the truth of the Hall–Paige conjecture on complete mappings of groups, whose proof depends on the classification of finite simple groups.

This has an application to the question of synchronization of finite automata, about which I spoke here in 2013.

It is conjectured that, if T has non-trivial cyclic Sylow 2-subgroups, then the Latin square graph of its Cayley table has chromatic number $|T| + 2$. We conjecture that the same is true for the diagonal graph for any even m.

An extension

I will briefly mention some work on extending this result.

An extension

I will briefly mention some work on extending this result. The main theorem talks about the situation where we have $m + 1$ partitions, any m of which are minimal elements of an m-dimensional Cartesian lattice.

An extension

I will briefly mention some work on extending this result. The main theorem talks about the situation where we have $m + 1$ partitions, any m of which are minimal elements of an m-dimensional Cartesian lattice.

Question

What if we have more than $m + 1$ partitions?

An extension

I will briefly mention some work on extending this result. The main theorem talks about the situation where we have $m + 1$ partitions, any m of which are minimal elements of an m-dimensional Cartesian lattice.

Question

What if we have more than $m + 1$ partitions?

For $m = 2$, we just have a set of mutually orthogonal Latin squares. So we called the general case a set of mutually orthogonal diagonal semilattices, or MODS.

An extension

I will briefly mention some work on extending this result.
The main theorem talks about the situation where we have $m+1$ partitions, any m of which are minimal elements of an m-dimensional Cartesian lattice.

Question

What if we have more than $m+1$ partitions?

For $m = 2$, we just have a set of mutually orthogonal Latin squares. So we called the general case a set of mutually orthogonal diagonal semilattices, or MODS.
We have a little theory and some examples. But we can't even settle the following question. Any $m+1$ of the partitions define a diagonal semilattice, which is coordinatised by a group. Does every set of $m+1$ partitions define the same group?

An extension

I will briefly mention some work on extending this result.
The main theorem talks about the situation where we have $m + 1$ partitions, any m of which are minimal elements of an m-dimensional Cartesian lattice.

Question

What if we have more than $m + 1$ partitions?

For $m = 2$, we just have a set of mutually orthogonal Latin squares. So we called the general case a set of mutually orthogonal diagonal semilattices, or MODS.
We have a little theory and some examples. But we can't even settle the following question. Any $m + 1$ of the partitions define a diagonal semilattice, which is coordinatised by a group. Does every set of $m + 1$ partitions define the same group?
This is false for MOLS. We can have two orthogonal Latin squares (defining four partitions) such that any three of the four define a group, but only one pair of the groups are isomorphic.

The example

Here is the example. The four partitions are rows, columns, first letters, second letters.

The example

Here is the example. The four partitions are rows, columns, first letters, second letters.

11	22	33	44	55	66	77	88
42	34	21	13	86	78	65	57
53	61	74	82	17	25	38	46
84	73	62	51	48	37	26	15
35	47	16	28	71	83	52	64
76	85	58	67	32	41	14	23
27	18	45	36	63	54	81	72
68	56	87	75	24	12	43	31

Omitting the ith partition, for $i = 1, \ldots, 4$, we obtain the Cayley tables of the groups D_8, $C_2 \times C_4$, D_8, and $C_2 \times C_2 \times C_2$.

The example

Here is the example. The four partitions are rows, columns, first letters, second letters.

11	22	33	44	55	66	77	88
42	34	21	13	86	78	65	57
53	61	74	82	17	25	38	46
84	73	62	51	48	37	26	15
35	47	16	28	71	83	52	64
76	85	58	67	32	41	14	23
27	18	45	36	63	54	81	72
68	56	87	75	24	12	43	31

Omitting the ith partition, for $i = 1, \ldots, 4$, we obtain the Cayley tables of the groups D_8, $C_2 \times C_4$, D_8, and $C_2 \times C_2 \times C_2$.
Can we get four different groups in this way? What about more than four partitions?

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.
- R. A. Bailey, Peter J. Cameron, Cheryl E. Praeger, Csaba Schneider, The geometry of diagonal groups, *Trans. Amer. Math. Soc.*, in press; arXiv `2007.10726`
- R. A. Bailey, Peter J. Cameron, Michael Kinyon and Cheryl E. Praeger, Diagonal groups and arcs over groups, *Designs, Codes, Cryptography*, in press; arXiv `2010.16338`
- R. A. Bailey and Peter J. Cameron, The diagonal graph, *J. Ramanujan Math. Soc.* **36** (2021), 353–361. arXiv `2101.02451`

Thank you . . .

. . . for your attention.