Combinatorial properties of transformation monoids

Peter J. Cameron

35ACCMCC
Melbourne, December 2011

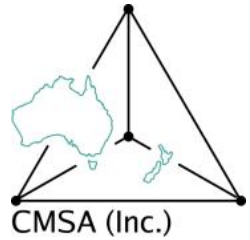

A few combinatorial problems . . .

- Is there an analogue of Baranyai's Theorem over fields with more than one element? That is, if k divides n, can one partition the set of all k-dimensional subspaces of an n-dimensional vector space into spreads, each spread containing every non-zero vector once?

A few combinatorial problems . . .

- Is there an analogue of Baranyai's Theorem over fields with more than one element? That is, if k divides n, can one partition the set of all k-dimensional subspaces of an n-dimensional vector space into spreads, each spread containing every non-zero vector once?
- Which polar spaces have ovoids, spreads, or partitions into ovoids?

A few combinatorial problems . . .

- Is there an analogue of Baranyai's Theorem over fields with more than one element? That is, if k divides n, can one partition the set of all k-dimensional subspaces of an n-dimensional vector space into spreads, each spread containing every non-zero vector once?
- Which polar spaces have ovoids, spreads, or partitions into ovoids?
- For which n can we partition the k-element subsets of an n-set into Steiner systems $S(3,4,n)$, or into Steiner systems $S(2,4,n)$?

. . . which all have something in common

I will define two properties of permutation groups, *synchronization* and *separation*. It turns out that "separating" implies "synchronizing", which implies "primitive" (and even "basic", in terms of the O'Nan–Scott classification).

. . . which all have something in common

I will define two properties of permutation groups, *synchronization* and *separation*. It turns out that "separating" implies "synchronizing", which implies "primitive" (and even "basic", in terms of the O'Nan–Scott classification).
But deciding which basic primitive groups are synchronizing, or separating, involves almost no group theory, and turns into a combinatorial problem, usually an interesting (and hard) problem. So this machine gives us a big supply of interesting combinatorial problems.

Automata

"Automaton" here means "finite deterministic automaton".

Automata

"Automaton" here means "finite deterministic automaton". An automaton is a device which can be in any one of a set Ω of internal states. On the console there are a number of coloured buttons; pressing a button forces the automaton to undergo a transition, a function from Ω to itself.

Automata

"Automaton" here means "finite deterministic automaton".
An automaton is a device which can be in any one of a set Ω of internal states. On the console there are a number of coloured buttons; pressing a button forces the automaton to undergo a transition, a function from Ω to itself.
Thus we can regard an automaton as an edge-coloured directed graph on Ω, with the property that there is a unique edge of each colour leaving each vertex.

Automata

"Automaton" here means "finite deterministic automaton".
An automaton is a device which can be in any one of a set Ω of internal states. On the console there are a number of coloured buttons; pressing a button forces the automaton to undergo a transition, a function from Ω to itself.
Thus we can regard an automaton as an edge-coloured directed graph on Ω, with the property that there is a unique edge of each colour leaving each vertex.
An automaton is synchronizing if there is a sequence of transitions which brings it into a fixed state $\alpha \in \Omega$, from any initial state.

The dungeon

You are in a dungeon consisting of a number of rooms. Passages are marked with coloured arrows. Each room contains a special door; in one room, the door leads to freedom, but in all the others, to instant death. You have a schematic map of the dungeon, but you do not know where you are.

The dungeon

You are in a dungeon consisting of a number of rooms. Passages are marked with coloured arrows. Each room contains a special door; in one room, the door leads to freedom, but in all the others, to instant death. You have a schematic map of the dungeon, but you do not know where you are.

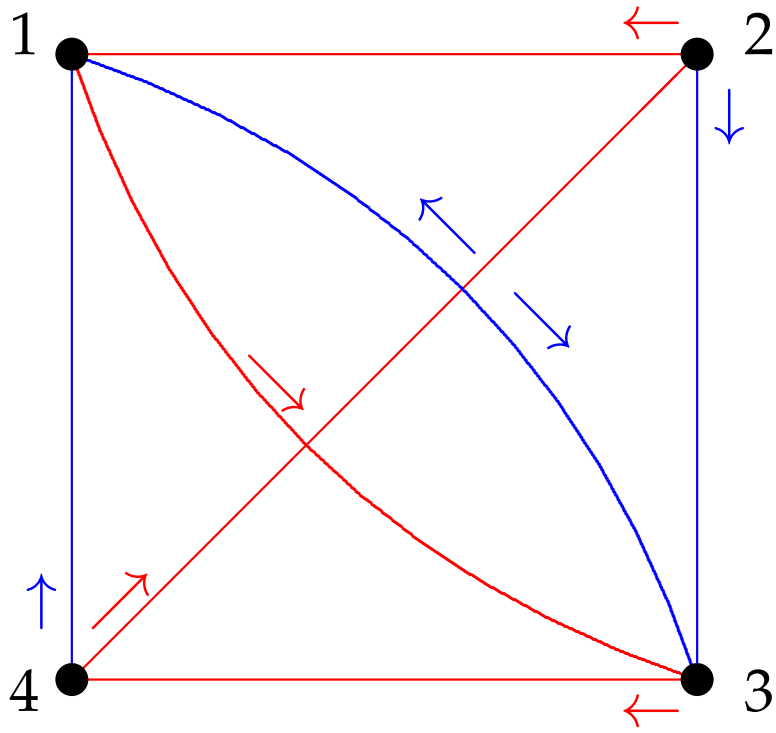

The dungeon

You are in a dungeon consisting of a number of rooms. Passages are marked with coloured arrows. Each room contains a special door; in one room, the door leads to freedom, but in all the others, to instant death. You have a schematic map of the dungeon, but you do not know where you are.

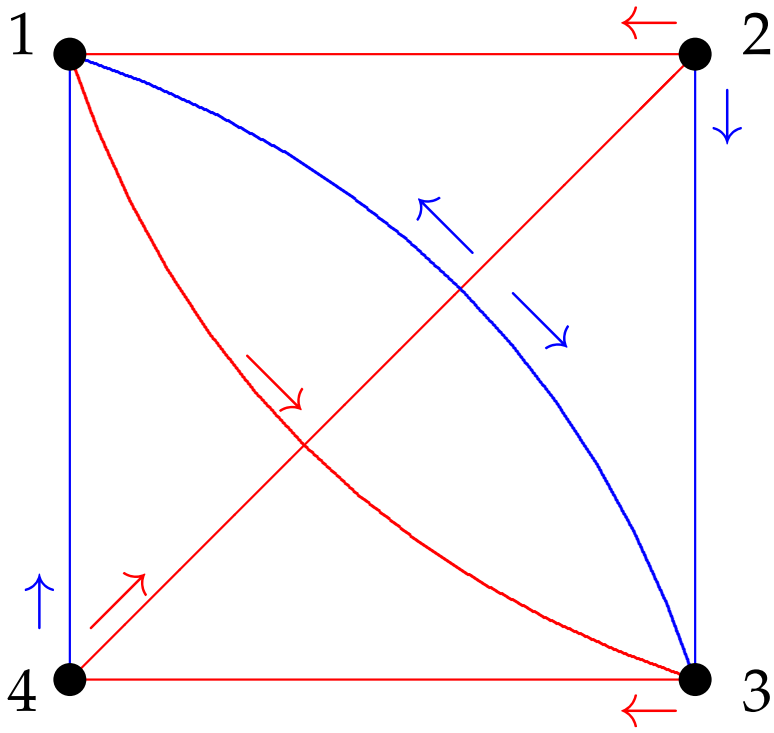

You can check that (Blue, Red, Blue, Blue) takes you to room 3 no matter where you start.

Algebraic formulation

Multiple button presses correspond to composition of transitions. The set of all functions generated by the given set S of transitions is closed under composition and contains the identity; thus it is a transformation monoid, the monoid generated by S.

Algebraic formulation

Multiple button presses correspond to composition of transitions. The set of all functions generated by the given set S of transitions is closed under composition and contains the identity; thus it is a transformation monoid, the monoid generated by S.
Note that any permutation in the monoid generated by S actually lies in the group generated by the permutations in S, since a product including a non-permutation cannot be a permutation.

Algebraic formulation

Multiple button presses correspond to composition of transitions. The set of all functions generated by the given set S of transitions is closed under composition and contains the identity; thus it is a transformation monoid, the monoid generated by S.
Note that any permutation in the monoid generated by S actually lies in the group generated by the permutations in S, since a product including a non-permutation cannot be a permutation.
An automaton is synchronizing if and only if this monoid contains a constant function (an element of rank 1). A word in the generators (that is, a series of button presses which evaluates to a constant function) is called a reset word.

The Černý conjecture

The study of synchronizing automata has been driven by the Černý conjecture, made in the 1960s and still open:

> *If an n-state automaton is synchronizing, then it has a reset word of length at most $(n-1)^2$.*

The Černý conjecture

The study of synchronizing automata has been driven by the Černý conjecture, made in the 1960s and still open:

> *If an n-state automaton is synchronizing, then it has a reset word of length at most $(n-1)^2$.*

If true, this would be best possible, as the following example shows.

An example

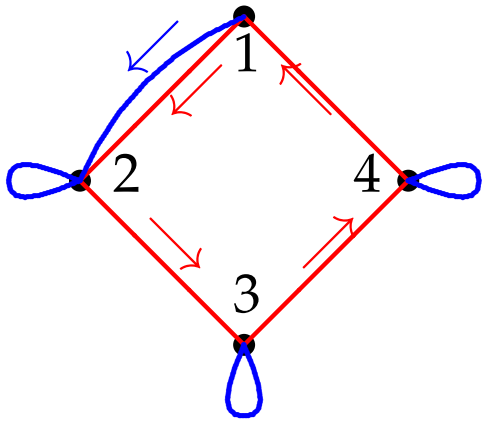

An example

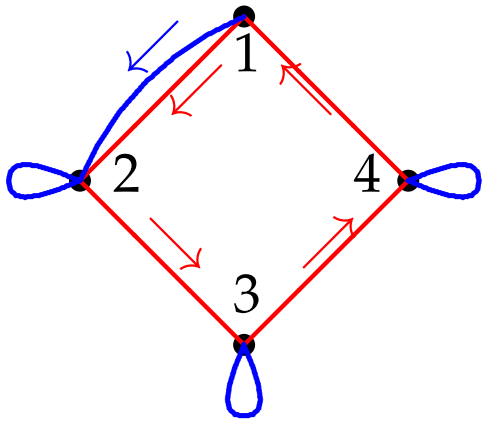

	B	R	R	R	B	R	R	R	B
1	2	3	4	1	2	3	4	1	2
2	2	3	4	1	2	3	4	1	2
3	3	4	1	2	2	3	4	1	2
4	4	1	2	3	3	4	1	2	2

An example

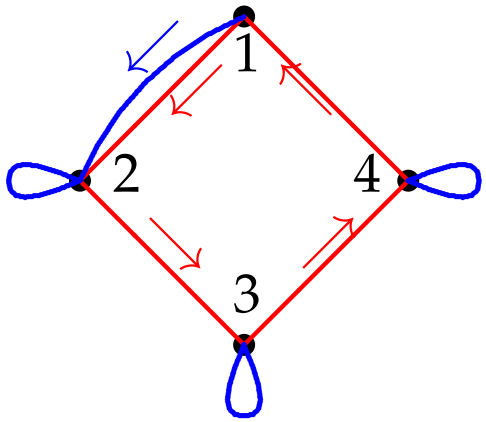

	B	R	R	R	B	R	R	R	B
1	2	3	4	1	2	3	4	1	2
2	2	3	4	1	2	3	4	1	2
3	3	4	1	2	2	3	4	1	2
4	4	1	2	3	3	4	1	2	2

So BRRRBRRRB is a reset word of length $9 = (4-1)^2$. It can be checked that this is the shortest reset word.

An example

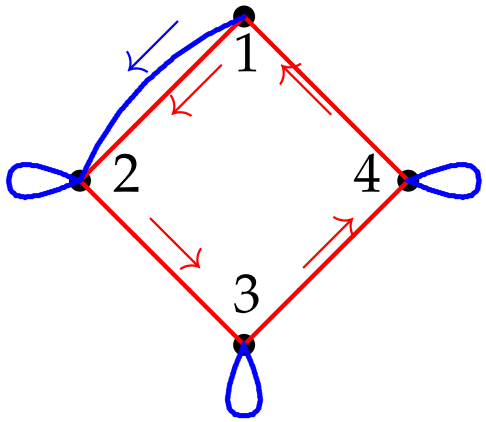

	B	R	R	R	B	R	R	R	B
1	2	3	4	1	2	3	4	1	2
2	2	3	4	1	2	3	4	1	2
3	3	4	1	2	2	3	4	1	2
4	4	1	2	3	3	4	1	2	2

So BRRRBRRRB is a reset word of length $9 = (4-1)^2$. It can be checked that this is the shortest reset word.

Replacing the square with a regular n-gon gives examples meeting the bound for all n.

Semigroups and groups

An approach to the Černý conjecture was introduced by João Araújo and Ben Steinberg.

Semigroups and groups

An approach to the Černý conjecture was introduced by João Araújo and Ben Steinberg.
We know that the permutations in a transformation monoid form a permutation group generated by the members of the generating set which are permutations. They made the following definition:

Semigroups and groups

An approach to the Černý conjecture was introduced by João Araújo and Ben Steinberg.
We know that the permutations in a transformation monoid form a permutation group generated by the members of the generating set which are permutations. They made the following definition:
Let G be a permutation group on Ω. We say that G is synchronizing if, whenever $f : \Omega \to \Omega$ is not a permutation, the monoid $\langle G, f \rangle$ generated by G and f contains an element of rank 1 (i.e. is synchronizing in the earlier sense).

Semigroups and groups

An approach to the Černý conjecture was introduced by João Araújo and Ben Steinberg.
We know that the permutations in a transformation monoid form a permutation group generated by the members of the generating set which are permutations. They made the following definition:
Let G be a permutation group on Ω. We say that G is synchronizing if, whenever $f : \Omega \to \Omega$ is not a permutation, the monoid $\langle G, f \rangle$ generated by G and f contains an element of rank 1 (i.e. is synchronizing in the earlier sense).
The hope was that, if G is synchronizing, we can use its structure to bound the length of the reset word in $\langle G, f \rangle$.

Synchronizing groups

The preceding synchronizing monoid has two generators, one a permutation which generates the cyclic group of order 4.

Synchronizing groups

The preceding synchronizing monoid has two generators, one a permutation which generates the cyclic group of order 4.
But C_4 is not a synchronizing group. If its generator is $g = (1, 2, 3, 4)$, and if f maps $1 \mapsto 1, 2 \mapsto 2, 3 \mapsto 1$ and $4 \mapsto 2$, then it is easy to see that any word containing at least one occurrence of f has an edge of the 4-cycle as its image.

Synchronizing implies primitive

A permutation group G on Ω is primitive if it preserves no equivalence relation on G except for the two trivial ones (equality and the "universal" relation.

Synchronizing implies primitive

A permutation group G on Ω is primitive if it preserves no equivalence relation on G except for the two trivial ones (equality and the "universal" relation.
Now a synchronizing group is primitive. For suppose that G is imprimitive. Choose a transversal T for the non-trivial G-invariant equivalence relation, and let f map each point to its representative in T. Then any word involving f and elements of G evaluates to a function whose image is a transversal to the equivalence relation.

Another characterisation of synchronization

João Araújo showed:

> *A permutation group G on Ω is non-synchronizing if and only if there exists a non-trivial partition π of Ω and a subset T of Ω such that every image of T under G is a transversal of π. It is synchronizing if no such pair exists.*

Another characterisation of synchronization

João Araújo showed:

> *A permutation group G on Ω is non-synchronizing if and only if there exists a non-trivial partition π of Ω and a subset T of Ω such that every image of T under G is a transversal of π. It is synchronizing if no such pair exists.*

The proof is almost identical to the argument just given.

Synchronizing implies basic

The celebrated O'Nan–Scott Theorem classifies primitive groups into a number of types. We can exclude one of these types for synchronizing groups.

Synchronizing implies basic

The celebrated O'Nan–Scott Theorem classifies primitive groups into a number of types. We can exclude one of these types for synchronizing groups.

A permutation group G on Ω is *non-basic* if it preserves a Cartesian power (hypercube) structure on Ω, and is *basic* otherwise.

Synchronizing implies basic

The celebrated O'Nan–Scott Theorem classifies primitive groups into a number of types. We can exclude one of these types for synchronizing groups.

A permutation group G on Ω is *non-basic* if it preserves a Cartesian power (hypercube) structure on Ω, and is *basic* otherwise.

Now a synchronizing group is basic. This is illustrated in the picture, where T is simultaneously a transversal for the partitions of Ω into each parallel class of codimension-1 faces.

Synchronizing implies basic

The celebrated O'Nan–Scott Theorem classifies primitive groups into a number of types. We can exclude one of these types for synchronizing groups.

A permutation group G on Ω is *non-basic* if it preserves a Cartesian power (hypercube) structure on Ω, and is *basic* otherwise.

Now a synchronizing group is basic. This is illustrated in the picture, where T is simultaneously a transversal for the partitions of Ω into each parallel class of codimension-1 faces.

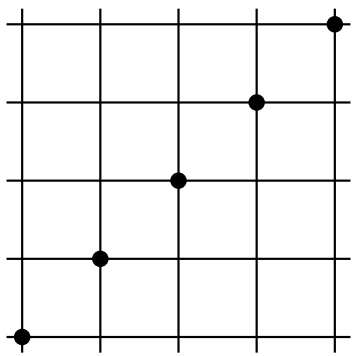

Separating groups

This class of groups is not directly inspired by automata theory but has a close relationship to synchronizing groups.

Separating groups

This class of groups is not directly inspired by automata theory but has a close relationship to synchronizing groups.
Let G act transitively on Ω. We say that G is separating if there do not exist subsets A, B of Ω such that $|A|, |B| > 1$ and $|Ag \cap B| = 1$ for all $g \in G$.

Separating groups

This class of groups is not directly inspired by automata theory but has a close relationship to synchronizing groups.
Let G act transitively on Ω. We say that G is separating if there do not exist subsets A, B of Ω such that $|A|, |B| > 1$ and $|Ag \cap B| = 1$ for all $g \in G$.
A separating group is synchronizing. For if G is not synchronizing, let π be a non-trivial partition and T a subset such that Tg is a transversal to π for all $g \in G$; taking $A = T$ and B a part of π shows that G is not separating.

Separating groups

This class of groups is not directly inspired by automata theory but has a close relationship to synchronizing groups.
Let G act transitively on Ω. We say that G is separating if there do not exist subsets A, B of Ω such that $|A|, |B| > 1$ and $|Ag \cap B| = 1$ for all $g \in G$.
A separating group is synchronizing. For if G is not synchronizing, let π be a non-trivial partition and T a subset such that Tg is a transversal to π for all $g \in G$; taking $A = T$ and B a part of π shows that G is not separating.
It is easy to see (the proof comes later) that a 2-transitive group is separating. So we have a hierarchy:

$$\begin{aligned} \text{2-transitive} &\Rightarrow \text{separating} \Rightarrow \text{synchronizing} \\ &\Rightarrow \text{basic} \Rightarrow \text{primitive.} \end{aligned}$$

A characterisation

There are polynomial-time algorithms to test whether a permutation group is transitive, 2-transitive or primitive.

A characterisation

There are polynomial-time algorithms to test whether a permutation group is transitive, 2-transitive or primitive. Such tests are not known for synchronization or separation. But there are tests which can be applied in practice for groups with degrees into the thousands. They depend on the following characterisations, most conveniently stated in negative form.

A characterisation

There are polynomial-time algorithms to test whether a permutation group is transitive, 2-transitive or primitive. Such tests are not known for synchronization or separation. But there are tests which can be applied in practice for groups with degrees into the thousands. They depend on the following characterisations, most conveniently stated in negative form. A graph is non-trivial if it is not complete or null. ω, α and χ denote the clique number, independence number, and chromatic number of a graph.

Theorem

- *The permutation group G on Ω is non-synchronizing if and only if there is a non-trivial G-invariant graph X with $\chi(X) = \omega(X)$.*

Theorem

- *The permutation group G on Ω is non-synchronizing if and only if there is a non-trivial G-invariant graph X with $\chi(X) = \omega(X)$.*
- *The transitive permutation group G on Ω is non-separating if and only if there is a non-trivial G-invariant graph X with $\omega(X) \cdot \alpha(X) = |\Omega|$.*

Theorem

- *The permutation group G on Ω is non-synchronizing if and only if there is a non-trivial G-invariant graph X with $\chi(X) = \omega(X)$.*
- *The transitive permutation group G on Ω is non-separating if and only if there is a non-trivial G-invariant graph X with $\omega(X) \cdot \alpha(X) = |\Omega|$.*

Note that the second condition is the same for a graph and its complement.

An algorithm

Thus we have the following algorithm to check whether the group G is synchronizing or separating:

An algorithm

Thus we have the following algorithm to check whether the group G is synchronizing or separating:

- Construct all the non-trivial G-invariant graphs (there are $2^d - 2$ of these, falling into $2^{d-1} - 1$ complementary pairs), where d is the number of G-orbits on unordered pairs from Ω.

An algorithm

Thus we have the following algorithm to check whether the group G is synchronizing or separating:

- Construct all the non-trivial G-invariant graphs (there are $2^d - 2$ of these, falling into $2^{d-1} - 1$ complementary pairs), where d is the number of G-orbits on unordered pairs from Ω.
- For one graph X of each pair, check whether it satisfies $\omega(X) \cdot \alpha(X) = |\Omega|$. If no such graph exists, the group is separating (and hence synchronizing).

An algorithm

Thus we have the following algorithm to check whether the group G is synchronizing or separating:

- Construct all the non-trivial G-invariant graphs (there are $2^d - 2$ of these, falling into $2^{d-1} - 1$ complementary pairs), where d is the number of G-orbits on unordered pairs from Ω.
- For one graph X of each pair, check whether it satisfies $\omega(X) \cdot \alpha(X) = |\Omega|$. If no such graph exists, the group is separating (and hence synchronizing).
- For each of the graphs found in the previous step, and their complements, check whether $\omega(X) = \chi(X)$ (note that $\omega(X)$ is already known). If one is found, then G is not synchronizing; otherwise it is.

Efficiency

This looks like a very bad algorithm: we have exponentially many graphs, and have to calculate their clique numbers and possibly their chromatic numbers (these are NP-hard).

Efficiency

This looks like a very bad algorithm: we have exponentially many graphs, and have to calculate their clique numbers and possibly their chromatic numbers (these are NP-hard).
But it is not as bad as it looks.

Efficiency

This looks like a very bad algorithm: we have exponentially many graphs, and have to calculate their clique numbers and possibly their chromatic numbers (these are NP-hard).
But it is not as bad as it looks.

- In many important cases, $d = 2$, so we have only one graph to consider in step 2, and two in step 3.

Efficiency

This looks like a very bad algorithm: we have exponentially many graphs, and have to calculate their clique numbers and possibly their chromatic numbers (these are NP-hard).
But it is not as bad as it looks.

- In many important cases, $d = 2$, so we have only one graph to consider in step 2, and two in step 3.
- Finding clique number (especially) of highly symmetric graphs is possible for quite large graphs, using algorithms that exploit the symmetry. (This is the philosophy of the GAP share package Grape.)

Efficiency

This looks like a very bad algorithm: we have exponentially many graphs, and have to calculate their clique numbers and possibly their chromatic numbers (these are NP-hard).
But it is not as bad as it looks.

- In many important cases, $d = 2$, so we have only one graph to consider in step 2, and two in step 3.
- Finding clique number (especially) of highly symmetric graphs is possible for quite large graphs, using algorithms that exploit the symmetry. (This is the philosophy of the GAP share package Grape.)

All primitive groups of degree up to 400, and interesting groups with degrees in the thousands, have been tested.

A class of problems

So we need to consider the general problem:

> *Let X be a graph on n vertices admitting a primitive basic automorphism group. Decide whether $\omega(X) = \chi(X)$, and whether $\omega(X) \cdot \alpha(X) = n$.*

A class of problems

So we need to consider the general problem:

> *Let X be a graph on n vertices admitting a primitive basic automorphism group. Decide whether $\omega(X) = \chi(X)$, and whether $\omega(X) \cdot \alpha(X) = n$.*

It turns out quite often that maximal cliques and cocliques and minimal colourings of such graphs are of great combinatorial or geometric interest. I will give a number of examples.

S_n on pairs

We begin with a simple example, S_n acting on the set of 2-element subsets of $\{1, \ldots, n\}$. This group is primitive (and basic) if $n \geq 5$.

S_n on pairs

We begin with a simple example, S_n acting on the set of 2-element subsets of $\{1, \ldots, n\}$. This group is primitive (and basic) if $n \geq 5$.
There are just two orbits on pairs of two-element subsets, depending on the size of intersection of the two subsets (1 or 0). The corresponding pair of G-invariant graphs are $X = L(K_n)$, the line graph of K_n (otherwise known as the triangular graph) and its complement.

S_n on pairs

We begin with a simple example, S_n acting on the set of 2-element subsets of $\{1, \ldots, n\}$. This group is primitive (and basic) if $n \geq 5$.

There are just two orbits on pairs of two-element subsets, depending on the size of intersection of the two subsets (1 or 0). The corresponding pair of G-invariant graphs are $X = L(K_n)$, the line graph of K_n (otherwise known as the triangular graph) and its complement.

It is easy to see that $\omega(X) = n - 1$, $\alpha(X) = \lfloor n/2 \rfloor$. The product of these numbers is $\binom{n}{2}$ if and only if n is even. So, if n is odd, then G is separating, and hence synchronizing.

S_n on pairs

We begin with a simple example, S_n acting on the set of 2-element subsets of $\{1, \ldots, n\}$. This group is primitive (and basic) if $n \geq 5$.

There are just two orbits on pairs of two-element subsets, depending on the size of intersection of the two subsets (1 or 0). The corresponding pair of G-invariant graphs are $X = L(K_n)$, the line graph of K_n (otherwise known as the triangular graph) and its complement.

It is easy to see that $\omega(X) = n - 1$, $\alpha(X) = \lfloor n/2 \rfloor$. The product of these numbers is $\binom{n}{2}$ if and only if n is even. So, if n is odd, then G is separating, and hence synchronizing.

We have $\chi(X) = n - 1$, so $\omega(X) = \chi(X)$. But $\chi(\bar{X}) = n - 2$, which is greater than $n/2$ for $n \geq 6$. So, for even n, G is not synchronizing.

S_n on k-sets

The group $G = S_n$ acting on k-element subsets of $\{1, \ldots, n\}$ is primitive and basic for $n \geq 7$. However, its synchronizing properties are a little more complicated:

S_n on k-sets

The group $G = S_n$ acting on k-element subsets of $\{1, \ldots, n\}$ is primitive and basic for $n \geq 7$. However, its synchronizing properties are a little more complicated:

Theorem

For $n \geq 7$, the group S_n acting on 3-sets is synchronizing, or separating, if and only if $n \equiv 2, 4, 5 \pmod 6$ and $n \neq 8$.

S_n on k-sets

The group $G = S_n$ acting on k-element subsets of $\{1, \ldots, n\}$ is primitive and basic for $n \geq 7$. However, its synchronizing properties are a little more complicated:

Theorem

For $n \geq 7$, the group S_n acting on 3-sets is synchronizing, or separating, if and only if $n \equiv 2, 4, 5$ (mod 6) and $n \neq 8$.

I will sketch some of the ideas of the proof as an illustration of the combinatorics used.

The group G has three orbits on pairs of 3-sets, depending on the size (0, 1 or 2) of the intersection. So there are six non-trivial G-invariant graphs, falling into three complementary pairs. We denote by X_{12} the graph in which two triples are joined if they intersect in 1 or 2 points, and similarly for the others.

The group G has three orbits on pairs of 3-sets, depending on the size (0, 1 or 2) of the intersection. So there are six non-trivial G-invariant graphs, falling into three complementary pairs. We denote by X_{12} the graph in which two triples are joined if they intersect in 1 or 2 points, and similarly for the others.
Clearly $\omega(X_0) = \lfloor n/3 \rfloor = \alpha(X_{12})$. Baranyai's Theorem asserts that, if n is a multiple of 3, then the complete 3-hypergraph has a 1-factorisation; so $\chi(X_{12}) = \binom{n-1}{2}$. Thus G is not synchronizing if n is divisible by 3. Also, the Erdős–Ko–Rado theorem asserts that $\omega(X_{12}) = \binom{n-1}{2}$; so for n not divisible by 3, this pair of graphs does not need to be considered further.

Now $\omega(X_2) \leq n(n-1)/6$, with equality if and only if there exists a Steiner triple system of order n, that is, n is congruent to 1 or 3 mod 6. Also $\alpha(X_2) = n - 2$ for $n \geq 7$. Teirlinck's theorem asserts that, if n is an admissible order of a STS and $n > 7$, then the set of all triples can be partitioned into $n - 2$ Steiner triple systems; so for these values of n, G is not synchronizing, and for other values except 7, the graphs do not need to be considered further.

Now $\omega(X_2) \leq n(n-1)/6$, with equality if and only if there exists a Steiner triple system of order n, that is, n is congruent to 1 or 3 mod 6. Also $\alpha(X_2) = n - 2$ for $n \geq 7$. Teirlinck's theorem asserts that, if n is an admissible order of a STS and $n > 7$, then the set of all triples can be partitioned into $n - 2$ Steiner triple systems; so for these values of n, G is not synchronizing, and for other values except 7, the graphs do not need to be considered further.

Analysis of the third pair of graphs, with special attention to the cases $n = 7$ and $n = 8$, finishes the proof.

Now $\omega(X_2) \leq n(n-1)/6$, with equality if and only if there exists a Steiner triple system of order n, that is, n is congruent to 1 or 3 mod 6. Also $\alpha(X_2) = n - 2$ for $n \geq 7$. Teirlinck's theorem asserts that, if n is an admissible order of a STS and $n > 7$, then the set of all triples can be partitioned into $n - 2$ Steiner triple systems; so for these values of n, G is not synchronizing, and for other values except 7, the graphs do not need to be considered further.

Analysis of the third pair of graphs, with special attention to the cases $n = 7$ and $n = 8$, finishes the proof.

Things get more complicated for $k \geq 4$.

Projective groups on subspaces

Similar things occur here. For the projective groups acting on lines, we have to consider whether there is a parallelism of lines: a maximal clique consists of all lines through a point if the dimension is large enough.

Projective groups on subspaces

Similar things occur here. For the projective groups acting on lines, we have to consider whether there is a parallelism of lines: a maximal clique consists of all lines through a point if the dimension is large enough.
Again, for subspaces of higher dimension, things get more complicated, and we need projective analogues of the Erdős–Ko–Rado theorem (conjectured but not yet proved), Baranyai's theorem, and Teirlinck's theorem; in the last case we are in a desperate situation since we do not even know whether the projective analogues of Steiner triple systems exist!

Classical groups

Classical groups form a very important class of primitive basic groups, acting on the points of the corresponding polar spaces (that is, the points of projective space which are isotropic with respect to a bilinear or Hermitian form, or singular with respect to a quadratic form).

Classical groups

Classical groups form a very important class of primitive basic groups, acting on the points of the corresponding polar spaces (that is, the points of projective space which are isotropic with respect to a bilinear or Hermitian form, or singular with respect to a quadratic form).

An ovoid is a set of points meeting every maximal subspace of the polar space in one point, while a spread is a partition of the points into maximal subspaces.

Classical groups

Classical groups form a very important class of primitive basic groups, acting on the points of the corresponding polar spaces (that is, the points of projective space which are isotropic with respect to a bilinear or Hermitian form, or singular with respect to a quadratic form).

An ovoid is a set of points meeting every maximal subspace of the polar space in one point, while a spread is a partition of the points into maximal subspaces.

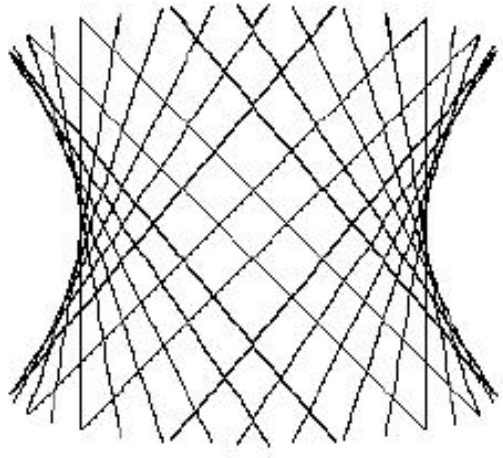

A quadric with two spreads

The classical group of a polar space has just two orbits on pairs of points (points may be perpendicular or not), and we find:

The classical group of a polar space has just two orbits on pairs of points (points may be perpendicular or not), and we find:

Theorem

- *G is non-separating if and only if the polar space possesses an ovoid.*

The classical group of a polar space has just two orbits on pairs of points (points may be perpendicular or not), and we find:

Theorem

- *G is non-separating if and only if the polar space possesses an ovoid.*
- *G is non-sychronizing if and only if the polar space possesses either an ovoid and a spread, or a partition into ovoids.*

The determination of which polar spaces have ovoids and spreads is far from complete, despite several decades of work by finite geometers. The existence of partitions into ovoids is a relatively new research problem, which has been resolved only in a few cases (if it is not trivially ruled out by the nonexistence of ovoids).

The determination of which polar spaces have ovoids and spreads is far from complete, despite several decades of work by finite geometers. The existence of partitions into ovoids is a relatively new research problem, which has been resolved only in a few cases (if it is not trivially ruled out by the nonexistence of ovoids).

This class gives our first examples of groups which are synchronizing but not separating. The polar spaces of the 5-dimensional orthogonal groups over finite fields of odd order have ovoids but no spreads, and Ball *et al.* have recently shown that they have no partitions into ovoids.

Further work

We have worked on some further classes of primitive groups, but there is a lot still to be done.

Further work

We have worked on some further classes of primitive groups, but there is a lot still to be done.
It seems likely that any class of primitive basic groups will throw up combinatorial problems comparable in difficulty, if not in interest, to those we have already met.

Further work

We have worked on some further classes of primitive groups, but there is a lot still to be done.
It seems likely that any class of primitive basic groups will throw up combinatorial problems comparable in difficulty, if not in interest, to those we have already met.
You are warmly invited to join the synchronization project!

Further work

We have worked on some further classes of primitive groups, but there is a lot still to be done.
It seems likely that any class of primitive basic groups will throw up combinatorial problems comparable in difficulty, if not in interest, to those we have already met.
You are warmly invited to join the synchronization project!
More information can be found in lecture notes from my short course:

`http://www.maths.qmul.ac.uk/~pjc/LTCC-2010-intensive3/`

Do two random elements synchronize?

This is motivated by a question of Brendan McKay.

Do two random elements synchronize?

This is motivated by a question of Brendan McKay.
Dixon's Theorem asserts:

Theorem
The probability that two random permutations of $\{1, 2, \ldots, n\}$ *generate the symmetric or alternating group tends to* 1 *as* $n \to \infty$.

Do two random elements synchronize?

This is motivated by a question of Brendan McKay.
Dixon's Theorem asserts:

Theorem
The probability that two random permutations of $\{1, 2, \ldots, n\}$ *generate the symmetric or alternating group tends to* 1 *as* $n \to \infty$.

We have to allow the alternating group since the probability that two random permutations are both even is $1/4$.

We cannot generate the full transformation monoid T_n with two elements, since we must include at least two permutations in any generating set. Moreover, permutations make up an exponentially small fraction of T_n. So we require many random elements to generate T_n with high probability.

We cannot generate the full transformation monoid T_n with two elements, since we must include at least two permutations in any generating set. Moreover, permutations make up an exponentially small fraction of T_n. So we require many random elements to generate T_n with high probability.
Instead, I make the following conjecture:

We cannot generate the full transformation monoid T_n with two elements, since we must include at least two permutations in any generating set. Moreover, permutations make up an exponentially small fraction of T_n. So we require many random elements to generate T_n with high probability.
Instead, I make the following conjecture:

Conjecture

The probability that two random elements of T_n generate a synchronizing monoid tends to 1 as $n \to \infty$.

Here is some data produced by James Mitchell. The first row is the number n, the second is the number of such pairs of elements of T_n generating a synchronizing monoid, the third is the total number n^{2n} of pairs of elements of T_n, and the fourth is the second divided by the third.

3	4	5	6
549	51520	8063385	1871446896
729	65536	9765625	2176782336
0.7531	0.7861	0.8257	0.8597

These results were obtained using the Citrus and Orb packages for GAP.

To prove this conjecture, following the proof of Dixon's Theorem, there are two steps:

To prove this conjecture, following the proof of Dixon's Theorem, there are two steps:

- Describe the maximal non-synchronizing submonoids of T_n;

To prove this conjecture, following the proof of Dixon's Theorem, there are two steps:

- Describe the maximal non-synchronizing submonoids of T_n;
- Use Inclusion-Exclusion to count the number of pairs of elements trapped in one of these submonoids, and show that it is $o(n^{2n})$.

To prove this conjecture, following the proof of Dixon's Theorem, there are two steps:

- Describe the maximal non-synchronizing submonoids of T_n;
- Use Inclusion-Exclusion to count the number of pairs of elements trapped in one of these submonoids, and show that it is $o(n^{2n})$.

The first step has been achieved: the maximal non-sychronizing monoids have been characterised in terms of graphs, though there is still a gap between the necessary and sufficient conditions. Certainly, we do not understand these submonoids well enough to take the second step.

Monoids and graphs

There is a connection between transformation monoids and graphs, with some features of a Galois correspondence. We define maps in each direction between transformation monoids on Ω and graphs on the vertex set Ω.

Monoids and graphs

There is a connection between transformation monoids and graphs, with some features of a Galois correspondence. We define maps in each direction between transformation monoids on Ω and graphs on the vertex set Ω.
Given a graph X, an endomorphism of X is a function on Ω which maps edges of X to edges. (We do not care what it does to non-edges, which may be mapped to non-edges or to edges or to single vertices). The endomorphisms of X clearly form a monoid $\mathrm{End}(X)$.

Monoids and graphs

There is a connection between transformation monoids and graphs, with some features of a Galois correspondence. We define maps in each direction between transformation monoids on Ω and graphs on the vertex set Ω.

Given a graph X, an endomorphism of X is a function on Ω which maps edges of X to edges. (We do not care what it does to non-edges, which may be mapped to non-edges or to edges or to single vertices). The endomorphisms of X clearly form a monoid $\mathrm{End}(X)$.

In the other direction, given a transformation monoid M, we define a graph $X = \mathrm{Gr}(M)$ by the rule that two vertices v, w are adjacent if and only if there does not exist $f \in M$ such that $vf = wf$.

Theorem

For any transformation monoid M, $\omega(\mathrm{Gr}(M)) = \chi(\mathrm{Gr}(M))$, and this number is equal to the minimum rank of an element of M.

Theorem

For any transformation monoid M, $\omega(\mathrm{Gr}(M)) = \chi(\mathrm{Gr}(M))$, and this number is equal to the minimum rank of an element of M.

Hence $\mathrm{Gr}(M)$ is complete if and only if M is a permutation group; and $\mathrm{Gr}(M)$ is null if and only if M is synchronizing.

Combining the maps

Theorem

For any transformation monoid M,

- $M \leq \mathrm{End}(\mathrm{Gr}(M))$;
- $\mathrm{Gr}(\mathrm{End}(\mathrm{Gr}(M))) = \mathrm{Gr}(M)$.

Combining the maps

Theorem

For any transformation monoid M,

- $M \leq \mathrm{End}(\mathrm{Gr}(M))$;
- $\mathrm{Gr}(\mathrm{End}(\mathrm{Gr}(M))) = \mathrm{Gr}(M)$.

The graph $\mathrm{Gr}(\mathrm{End}(X))$ is called the hull of X. We have $\mathrm{Hull}(\mathrm{Hull}(X)) = \mathrm{Hull}(X)$. In other words, a graph X is a hull if and only if it is its own hull (that is, $\mathrm{Hull}(X) = X$).

Another construction

Let X be a graph on the vertex set Ω with $\omega(X) = m$. Let X' be the spanning subgraph of X which consists of those edges of X which are contained in cliques of size m. Then
$\mathrm{End}(X) \leq \mathrm{End}(X')$.

Another construction

Let X be a graph on the vertex set Ω with $\omega(X) = m$. Let X' be the spanning subgraph of X which consists of those edges of X which are contained in cliques of size m. Then
$\mathrm{End}(X) \leq \mathrm{End}(X')$.
I will call Y the derived graph of X.

Maximal non-synchronizing monoids

Theorem

Let M be a maximal non-synchronizing submonoid of T_n. Then there are graphs X and Y on the vertex set $\Omega = \{1, \ldots, n\}$ satisfying the following conditions:

- $\mathrm{End}(X) = \mathrm{End}(Y) = M$;
- $\omega(X) = \omega(Y) = \chi(X) = \chi(Y)$;
- $X = \mathrm{Hull}(Y)$;
- $Y = X'$.

Maximal non-synchronizing monoids

Theorem

Let M be a maximal non-synchronizing submonoid of T_n. Then there are graphs X and Y on the vertex set $\Omega = \{1, \ldots, n\}$ satisfying the following conditions:

- $\mathrm{End}(X) = \mathrm{End}(Y) = M$;
- $\omega(X) = \omega(Y) = \chi(X) = \chi(Y)$;
- $X = \mathrm{Hull}(Y)$;
- $Y = X'$.

I do not know any examples where X and Y are not equal. If they are equal, then the converse holds:

Maximal non-synchronizing monoids

Theorem

Let M be a maximal non-synchronizing submonoid of T_n. Then there are graphs X and Y on the vertex set $\Omega = \{1, \ldots, n\}$ satisfying the following conditions:

- $\mathrm{End}(X) = \mathrm{End}(Y) = M$;
- $\omega(X) = \omega(Y) = \chi(X) = \chi(Y)$;
- $X = \mathrm{Hull}(Y)$;
- $Y = X'$.

I do not know any examples where X and Y are not equal. If they are equal, then the converse holds:

Theorem

Let X be a hull (other than the null graph), in which every edge is contained in a clique of size $\omega(X)$. Then $\mathrm{End}(X)$ is a maximal non-synchronizing submonoid of $T(\Omega)$.

There are many graphs satisfying the hypotheses of this theorem. The smallest consists of a single edge; there are $n(n-1)/2$ graphs of this form and each has $2n^{n-2}$ endomorphisms. So the probability that a random pair of endofunctions are both endomorphisms of a graph of this form is at most

$$\frac{n(n-1)}{2}\,\frac{n^{2(n-2)}}{n^{2n}} = O(n^{-2}).$$

There are many graphs satisfying the hypotheses of this theorem. The smallest consists of a single edge; there are $n(n-1)/2$ graphs of this form and each has $2n^{n-2}$ endomorphisms. So the probability that a random pair of endofunctions are both endomorphisms of a graph of this form is at most

$$\frac{n(n-1)}{2}\,\frac{n^{2(n-2)}}{n^{2n}} = O(n^{-2}).$$

This suggests that the probability that two random endofunctions generate a synchronizing monoid is at least $1 - O(1/n^2)$. However, we are still some way from a proof, since there are many graphs that need to be considered. Of course, there are big overlaps between their endomorphism monoids, so careful inclusion-exclusion is required!

There are many graphs satisfying the hypotheses of this theorem. The smallest consists of a single edge; there are $n(n-1)/2$ graphs of this form and each has $2n^{n-2}$ endomorphisms. So the probability that a random pair of endofunctions are both endomorphisms of a graph of this form is at most

$$\frac{n(n-1)}{2}\,\frac{n^{2(n-2)}}{n^{2n}} = O(n^{-2}).$$

This suggests that the probability that two random endofunctions generate a synchronizing monoid is at least $1 - O(1/n^2)$. However, we are still some way from a proof, since there are many graphs that need to be considered. Of course, there are big overlaps between their endomorphism monoids, so careful inclusion-exclusion is required!
For more details see `http://arxiv.org/abs/1108.3958`


End


THANK YOU