Graphs, groups and semigroups

Peter J. Cameron
University of St Andrews

Cracow graph theory conference
Rytro, September 2018

Terminology

Group theorists always begin a lecture with "Let G be a group", while graph theorists say "Let G be a graph". So we have a problem.
"Graph" is a Greek word, while "group" is a German word. So I should say "Let Γ be a graph" and "Let $\mathfrak{G}$ be a group".
However, after decades of lecturing, I never learned to write a Fraktur G on the blackboard, so my graphs will be Γ and my groups will be G.
In passing, notations change: in Richard Brauer's papers, $\mathfrak{G}$ is a group, its order is g, and a typical element is G. I find these papers quite hard to read!

Symmetry

As everyone knows, group theory is the part of mathematics that studies symmetry. But of course it is not quite so simple. Group theory and graph theory have been closely connected for many years. There are three areas where they have interacted:

- One area grows from Frucht's theorem: *Every group is the automorphism group of a graph*. This has been refined in many ways: we can ask for a graph with various properties, or for special types of group action.
- One area uses graphs just as a notation: the Bass-Serre theory of "graphs of groups", Coxeter groups, etc.
- And thirdly, finding the graphs with a rich group of automorphisms, and using graphs to study groups.

The first area is really about *restricting* symmetry, and the second is not really concerned with symmetry at all.
My talk will be about the third area.

Outline

I will start with a brief introduction to permutation groups, and talk about how in the 1960s graphs began to be used as a tool for studying permutation groups, in particular transitive (or primitive) groups.
Then I will say something about transformation semigroups, and give a much more recent connection between these and graphs. This will lead to a discussion of synchronization, a problem in automata theory which has led to some very challenging questions about graphs and groups.

Permutation groups

Permutation groups form the oldest part of group theory, featuring especially in the work of Galois (whom some regard as the founder of group theory).
Here is a brief primer on permutation group theory. The situation is that G is a permutation group on Ω, where $|\Omega| = n$; that is, G is a subgroup of the symmetric group S_n.
To make these definitions easier, I will say a structure on Ω is trivial if it is preserved by the whole symmetric group S_n, and non-trivial otherwise. Thus, for example,

- a subset of Ω is trivial if it is Ω or $\emptyset$, and non-trivial otherwise;
- a partition of Ω is trivial if and only if it is either the partition into singletons, or the partition with just one part;
- a graph on the vertex set Ω is trivial if and only if it is complete or null.

Permutation groups, 2

Let G be a permutation group on Ω.

- G is transitive if it fixes no non-trivial subset of Ω;
- G is primitive if it is transitive and fixes no non-trivial partition of Ω;
- G is 2-homogeneous if it fixes no non-trivial graph on Ω;
- G is 2-transitive if it fixes no non-trivial digraph on Ω.

These conditions form a hierarchy. We will add one more later. We have to require that a primitive group is transitive since otherwise the identity group acting on a set of two points would be primitive (there are only two partitions on such a set, both trivial).

Primitivity and graphs

There is a nice characterisation of primitivity in terms of graphs, due to Donald Higman.

Theorem

The permutation group G on Ω (with $|\Omega| > 2$) is primitive if and only if G fixes no non-null connected graph.

For, if G is imprimitive, it fixes a partition, and so preserves the disjoint union of complete graphs on the parts of the partition. Conversely, if G preserves a disconnected graph Γ, then it fixes the partition of Ω into connected components of Γ.

Orbits and transitivity

Given a permutation group G on Ω, we can define an equivalence relation on Ω by $x \equiv y$ if and only if there is an element $g \in G$ mapping x to y.
The reflexive, symmetric and transitive laws for $\equiv$ follow from the identity, inverse and closure axioms for a group.
The equivalence classes are the orbits of G; and G is transitive if there is just one orbit. Note that G fixes each orbit (as a set).
Also, G acts on the set $\Omega^{\{2\}}$ of 2-element subsets of Ω, and on the set Ω^2 of ordered pairs of elements. Now G is 2-homogeneous if and only if it is transitive on $\Omega^{\{2\}}$, and is 2-transitive if and only if its orbits on Ω^2 are the diagonal $\Delta = \{(x, x) : x \in \Omega\}$ and $\Omega^2 \setminus \Delta$.

A combinatorial tool

In the 1960s, Boris Weisfeiler in the USSR and Donald Higman in the USA independently considered the set of all orbits of G on Ω^2. These define a combinatorial object which Weisfeiler called a cellular algebra and Higman a coherent configuration; it generalises earlier ideas in group theory (Schur) and statistics (Bose).

Now the name cellular algebra has been hijacked for a different use, and so we call these objects coherent configurations.

If you want to learn more about these objects and their history, the slides from talks at a conference in Pilsen, Czech Republic, earlier this year, have a lot of information: see

`https://www.iti.zcu.cz/wl2018/slides.html`

Coherent configurations

Thus, a coherent configuration is a partition $\mathcal{P}$ of Ω^2 with the properties

- the diagonal is a union of parts of the partition;
- the transponse of a part of $\mathcal{P}$ (obtained by reversing all the ordered pairs) is a part of $\mathcal{P}$;
- if P_i, P_j, P_k are in $\mathcal{P}$, and $(x,y) \in P_k$, then the number a_{ij}^k of points $z \in \Omega$ such that $(x,z) \in P_i$ and $(z,y) \in P_j$ depends only on (i,j,k) and not on the choice of x, y.

Coherent configurations can be regarded as special edge-colourings of the complete directed graph.

The third condition

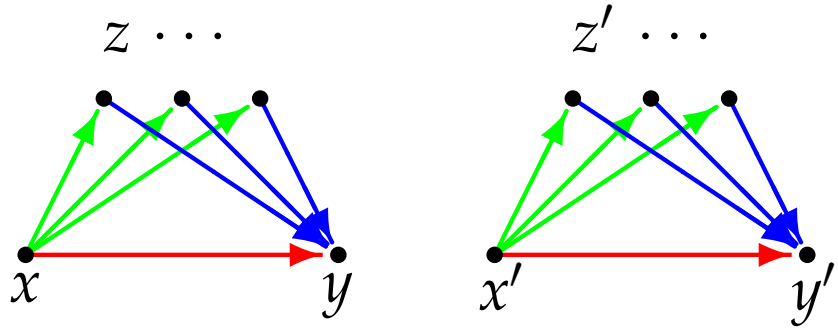

Colours denote orbits of G. An element $g \in G$ mapping (x, y) to (x', y') maps the set of points z shown to the set of points z'. So the numbers $a_{ij}{}^{k}$ of these points are independent of the choice of (x, y).

Algebra

Coherent configurations have a very nice algebraic structure. A set P of ordered pairs of elements of Ω can be represented by an $n \times n$ zero-one matrix A, whose (i,j) entry is 1 if $(i,j) \in P$ and 0 otherwise.
Let $\{A_1, \ldots, A_r\}$ be the matrices corresponding to the parts of the partition $\mathcal{P}$ of a coherent configuration. The axioms translate like this:

- the sum of all the A_i is the all-1 matrix J;
- there is a subset of the A_i summing to the identity matrix I;
- the set $\{A_1, \ldots, A_r\}$ is closed under transposition;
- for any i, j, we have $A_i A_j = \sum_{k=1}^{r} a_{ij}^k A_k$.

Thus the span of the matrices $A_1, \ldots, A_k$ is closed under multiplication, and so is a semisimple associative algebra.

Applications

By Wedderburn's theorem it is isomorphic to a direct sum of matrix algebras. The dimensions of these algebras and their multiplicities can in principle be calculated from the parameters a_{ij}^k; since these must be integers, we obtain necessary conditions on these parameters.
This is familiar in the case of strongly regular graphs. For example, the proof that a Moore graph (a graph with diameter 2 and girth 5) must have valency 2, 3, 7 or 57 uses this technique. For the first three values, there exists a unique and highly symmetric graph. The existence of a Moore graph of valency 57 is still unknown.

What are they good for?

Coherent configurations are central objects in algebraic combinatorics. They include many interesting combinatorial and geometric structures such as strongly regular or distance-regular graphs, projective planes, generalized polygons, symmetric and quasi-symmetric designs, Latin squares and sets of pairwise orthogonal Latin squares, Steiner triple systems, etc.
They also play an important role in László Babai's recent quasi-polynomial algorithm for graph isomorphism, the most important recent breakthrough on this topic. (They have always been connected with the graph isomorphism problem: this was Weisfeiler's original motivation.)

Graphs and monoids

We saw that there are many graphs associated with a permutation group; the coherent configuration is a way of packing them all into a single gadget.
By contrast, there is a natural construction of a graph from a transformation monoid, which happens to give the complete graph if applied to a permutation group.
I will discuss this, and its connection with synchronizing automata.

Endomorphisms

An endomorphism of a graph Γ is a map from the vertex set of Γ to itself which maps edges to edges. What happens to non-edges is not specified; a non-edge might map to a non-edge, or to an edge, or collapse to a single vertex.
For example, let $\{v, w\}$ be an edge of the bipartite graph Γ. The map that takes each vertex to the vertex of $\{v, w\}$ in the same bipartite block is an endomorphism.
The composition of endomorphisms is an endomorphism, and the identity map is an endomorphism. So the endomorphisms of a graph Γ form a transformation monoid $\mathrm{End}(\Gamma)$ on the vertex set of Γ.

The graph of a monoid

In the other direction, let M be a transformation monoid on Γ. Define a graph $\mathrm{Gr}(M)$ on Γ by the rule that $\{v, w\}$ is an edge if and only if there *does not* exist $t \in M$ such that $vt = wt$.

Theorem

- $\mathrm{End}(\mathrm{Gr}(M))) \geq M$.
- $\mathrm{Gr}(\mathrm{End}(\mathrm{Gr}(M))) = \mathrm{Gr}(M)$.
- *If* $M_1 \leq M_2$, *thten* $\mathrm{Gr}(M_1)$ *is a spanning subgraph of* $\mathrm{Gr}(M_2)$.
- $\mathrm{Gr}(M)$ *has clique number equal to chromatic number.*

All parts are easy. For example, let $t \in M$. If $\{x, y\}$ is an edge of $\mathrm{Gr}(M)$, but $\{xt, yt\}$ is not, then there exists $s \in M$ such that $xts = yts$, a contradiction. This proves the first statement.

Cores and hulls

The core of a graph Γ is the smallest induced subgraph Δ for which there is an endomorphism of Γ with image Δ. A graph may have many cores but they are all isomorphic (and all their endomorphisms are automorphisms).
A dual concept is the hull of Γ, which is simply $\mathrm{Gr}(\mathrm{End}(\Gamma))$. It contains Γ as a spanning subgraph. For example, if Γ is a path of length 3, then no endomorphism collapses the ends of Γ, so the hull of Γ is a 4-cycle. Note that taking the hull doesn't decrease the endomorphism monoid (or automorphism group). The vertex set of the core of Γ carries a complete subgraph of the hull of Γ, which is the core of the hull.

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 death. You have a map of the dungeon,
but you do not know where you are.
Can you escape? In other words, is there a sequence of colours
such that, if you use the doors in this sequence from any
starting point, you will end in a known place?

An example

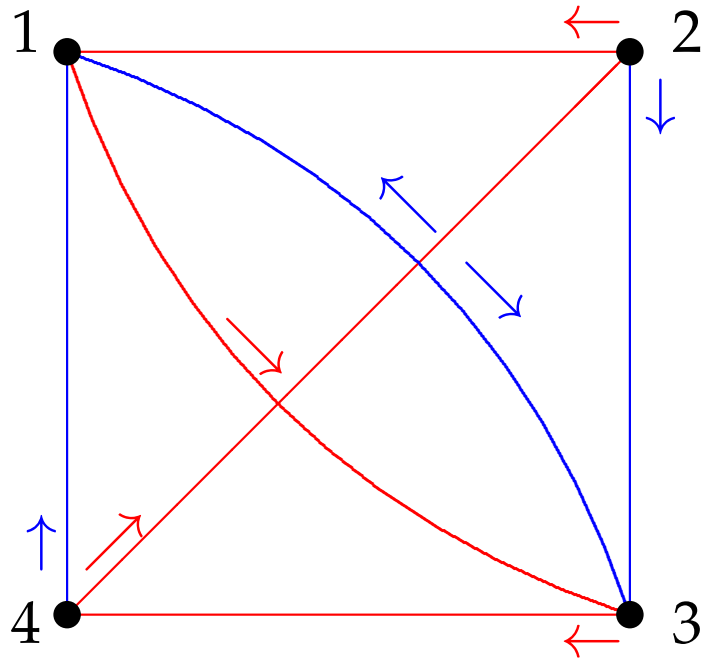

You can check that (Blue, Red, Blue) takes you to room 1 no matter where you start.

Automata

The diagram on the last page shows a finite-state deterministic automaton. This is a machine with a finite set of states, and a finite set of transitions, each transition being a map from the set of states to itself. The machine starts in an arbitrary state, and reads a word over an alphabet consisting of labels for the transitions (Red and Blue in the example); each time it reads a letter, it undergoes the corresponding transition.

A reset word is a word with the property that, if the automaton reads this word, it arrives at the same state, independent of its start state. An automaton which possesses a reset word is called synchronizing.

Not every finite automaton has a reset word. For example, if every transition is a permutation, then every word in the transitions evaluates to a permutation. How do we recognise when an automaton is synchronizing?

Industrial robotics

In a factory, parts are delivered by conveyor belt to a robot for assembly. Each part must be put on in the correct orientation. Assuming they arrive in random orientation, this is a job for a synchronizing automaton.
Suppose that the pieces are square, with a small projection on one side:

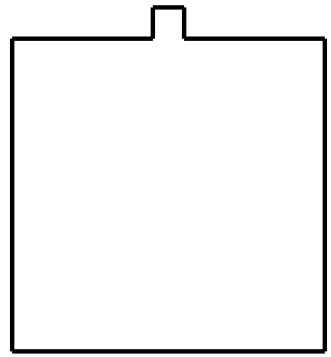

Suppose the conveyor has a square tray in which the pieces can lie in any orientation. Simple gadgets can be devised so that the first gadget rotates the square through 90° anticlockwise; the second rotates it only if it detects that the projection is pointing towards the top. The set-up can be represented by an automaton with four states and two transitions, see next slide.

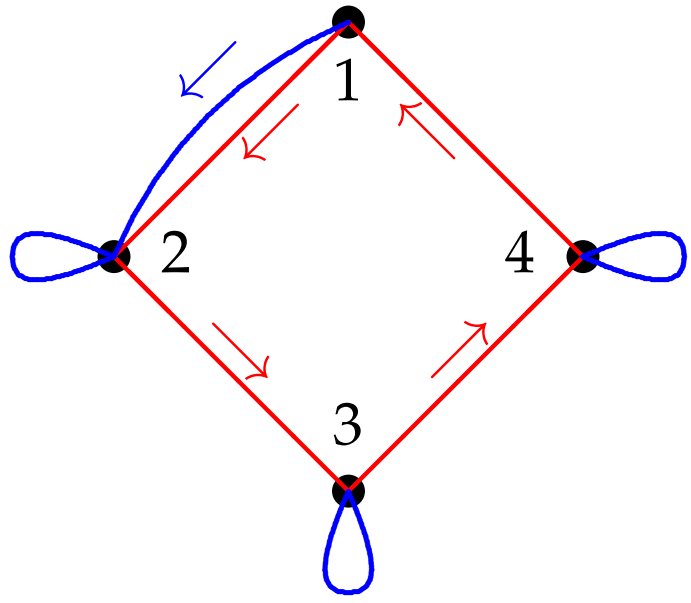

Now it can be verified that BRRRBRRRB is a reset word (and indeed that it is the shortest possible reset word for this automaton).

The Černý conjecture

This is a special case of the Černý conjecture, made about fifty years ago and still open:

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

The above example and the obvious generalisation show that the conjecture, if true, is best possible.
The Černý conjecture has been proved in some cases, but the best general upper bound known is $O(n^3)$, due to Pin. Here is a proof of an $O(n^3)$ bound, which does not get the best constant, but illustrates a simple but important principle.

Proposition

An automaton is synchronizing if and only if, for any two states a, b, there is a word in the transitions which takes the automaton to the same place starting from either a or b.

A bound

Now to obtain our bound, consider the diagram of the automaton extended to include pairs of states.

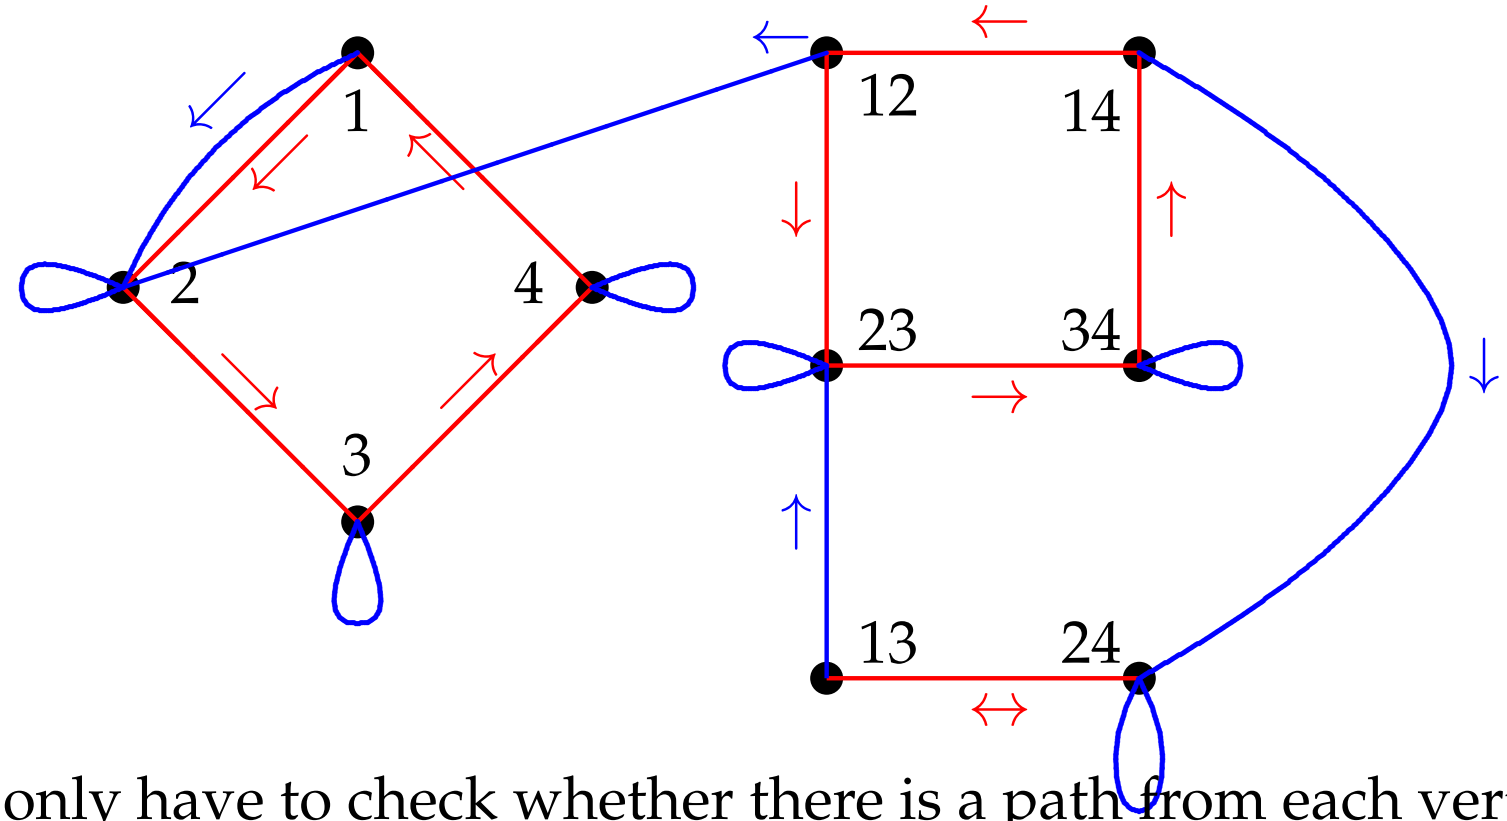

We only have to check whether there is a path from each vertex on the right (a pair of states) to a vertex on the left (a single state). Such a path (if it exists) has length $O(n^2)$, and we only require $n-1$ "collapses" of pairs to synchronize.

Automata and transformation semigroups

Combinatorially, an automaton is an edge-coloured digraph with one edge of each colour out of each vertex. Vertices are states, colours are transitions.
Algebraically, if $\Omega = \{1, \ldots, n\}$ is the set of states, then any transition is a map from Ω to itself. Reading a word composes the corresponding maps, so the set of maps corresponding to all words is a transformation semigroup on Ω.
So an automaton is a transformation semigroup with a distinguished generating set.

The obstruction to synchronization

From now on our graphs are simple and undirected. An endomorphism of a graph is a map from the graph to itself which takes edges to edges.
The endomorphisms of a graph Γ form a transformation semigroup; if Γ is not a null graph, then $\mathrm{End}(\Gamma)$ is not synchronizing, since edges cannot be collapsed.

Theorem

Let S be a transformation monoid on Ω. Then S fails to be synchronizing if and only if there exists a non-null graph Γ on the vertex set Ω for which $S \leq \mathrm{End}(\Gamma)$. Moreover, we may assume that $\omega(\Gamma) = \chi(\Gamma)$.

For the proof, simply take $\Gamma = \mathrm{Gr}(S)$.

Synchronizing groups

A permutation group is never synchronizing as a monoid, since no collapses at all occur.
We abuse language by making the following definition. A permutation group G on Ω is synchronizing if, for any map t on Ω which is not a permutation, the monoid $\langle G, t\rangle$ generated by G and t is synchronizing.

Theorem

A permutation group G on Ω is non-synchronizing if and only if there exists a G-invariant graph Γ, not complete or null, which has clique number equal to chromatic number.

Synchronization in the hierarchy

Recall that

$$\text{2-transitive} \Rightarrow \text{2-homogeneous} \Rightarrow \text{primitive} \Rightarrow \text{transitive}.$$

Theorem

- *A 2-homogeneous group is synchronizing.*
- *A synchronizing group is primitive.*

For a synchronizing group preserves no non-trivial graph with clique number equal to chromatic number, while a 2-homogeneous group preserves no non-trivial graph at all. On the other hand, an imprimitive group preserves a disjoint union of complete graphs of the same size; this certainlly has clique number equal to chromatic number.

How do you test?

There are polynomial-time algorithms for the other properties mentioned on the previous slide, but no such algorithm is known for synchronization.
The best we know is essentially the following. Given a permutation group G, which we may assume is primitive and not 2-homogeneous,

- Construct all the G-invariant graphs. There are $2^r - 2$ of these, where r is the number of orbits of G on 2-element subsets of Ω.
- Test whether these have clique number equal to chromatic number. If we find one that does, then G is non-synchronizing; otherwise it is synchronizing.

This involves potentially solving exponentially many NP-hard problems.

Problem

Can we do better?

Further reading:

- J. Araújo, P. J. Cameron and B. Steinberg, Between primitive and 2-transitive: synchronization and its friends, *Europ. Math. Soc. Surveys* **4** (2017), 101–184.
- Slides from *Symmetry vs Regularity*: `https://www.iti.zcu.cz/wl2018/slides.html`