Endomorphisms and synchronization

Peter J. Cameron

University of St Andrews

Algebraic Combinatorics – Godsil 65
Waterloo, June 2014

How to build a quantum computer

Happy birthday, Chris!

Around the time I started thinking about what I am talking about today, Chris was also thinking about it:
C. D. Godsil and G. F. Royle, Cores of geometric graphs, *Ann. Combinatorics* **15** (2011), 267–276.
I will explain why I am interested; I don't know why Chris and Gordon came to it ...
At the end of my talk I will mention one connection between our approaches.

Graphs and groups

The word "graph" has a Greek root, so a graph should be called Γ.
The word "group" has a Germanic origin, so a group should be $\mathfrak{G}$ (as it was for Richard Brauer).
I will break the second rule, since I don't believe that anyone not educated in Germany can recognise $\mathfrak{G}$ instantly. Sometimes I break the first rule too and call a graph X.

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.

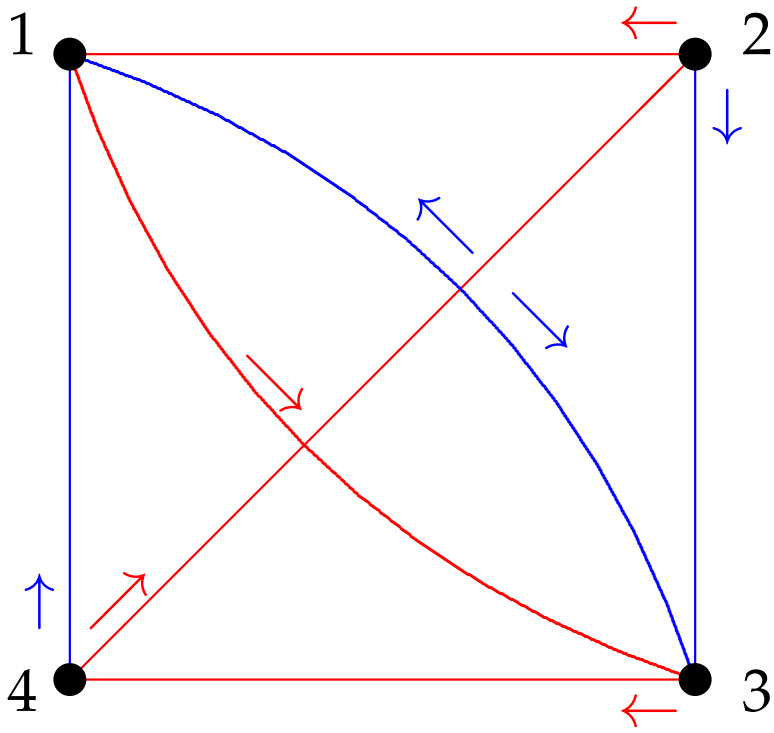

You can check that (Blue, Red, Blue, Blue) takes you to room 3 no matter where you start.

Definitions

A (finite deterministic) automaton consists of a finite set Ω of states, with a finite set S of transitions, maps from Ω to Ω.
The automaton is synchronizing if there is a word in the transitions which evaluates to a map of rank 1.
Combinatorially, an automaton is an edge-coloured directed graph on Ω such that every vertex is the source of a unique arc of each colour.
Algebraically, since we are interested in composing maps, an automaton is a transformation semigroup on Ω (a set of transformations closed under composition) with a prescribed set S of generators.

Quantum synchronization??

You will notice some correspondences between what follows and some of the talks on the "quantum afternoon" on Monday, and may wonder whether there is a definition of quantum synchronization.
However, from the point of view of physics, this makes no sense. Quantum evolution is always reversible, but synchronization is the ultimate irreversible process; when the automaton has synchronized, it has completely forgotten its past.

Another example

This example arises in industrial robotics.

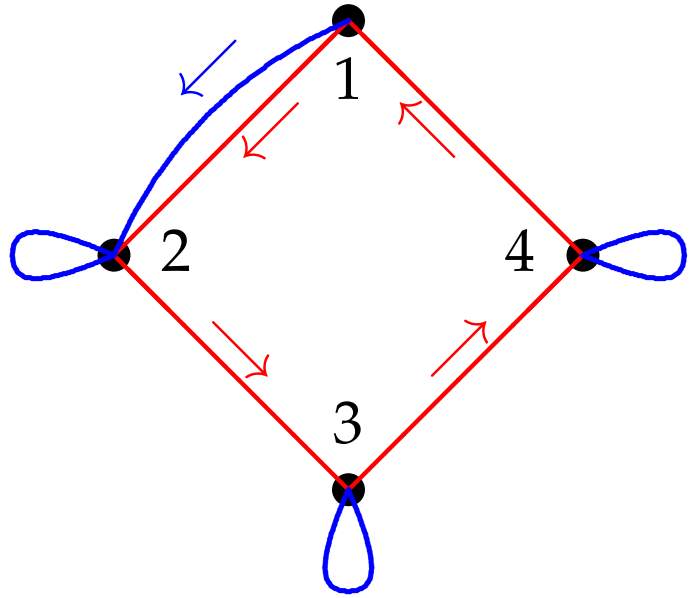

	B	R	R	R	B	R	R	R	B
1	2	3	4	1	2	3	4	1	2
2	2	3	4	1	2	3	4	1	2
3	3	4	1	2	2	3	4	1	2
4	4	1	2	3	3	4	1	2	2

So BRRRBRRRB is a reset word.

Problems

Problem (The Černý conjecture)

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

This problem is still open after nearly fifty years. The example on the previous slide and the obvious generalisation show that, if true, it is best possible.
Two related computational problems. Given an automaton (Ω, S),

- Decide whether it is synchronizing.
- If so, find the shortest reset word.

Testing synchronization

Proposition

An automaton (Ω, S) is synchronizing if and only if, for any two states $a, b \in \Omega$, there is a word $w_{a,b}$ in the elements of S which maps a and b to the same place.

Proof.

"Only if" is clear, so suppose that the condition holds. Let f be an element of $\langle S \rangle$ of smallest possible rank. If the rank of S is greater than 1, then choose two points a, b in the image; then fw_{ab} has smaller rank than f. So f has rank 1, and the automaton is synchronizing. □

So we only have to consider all pairs of states.

The picture shows the previous example, extended to pairs of states.

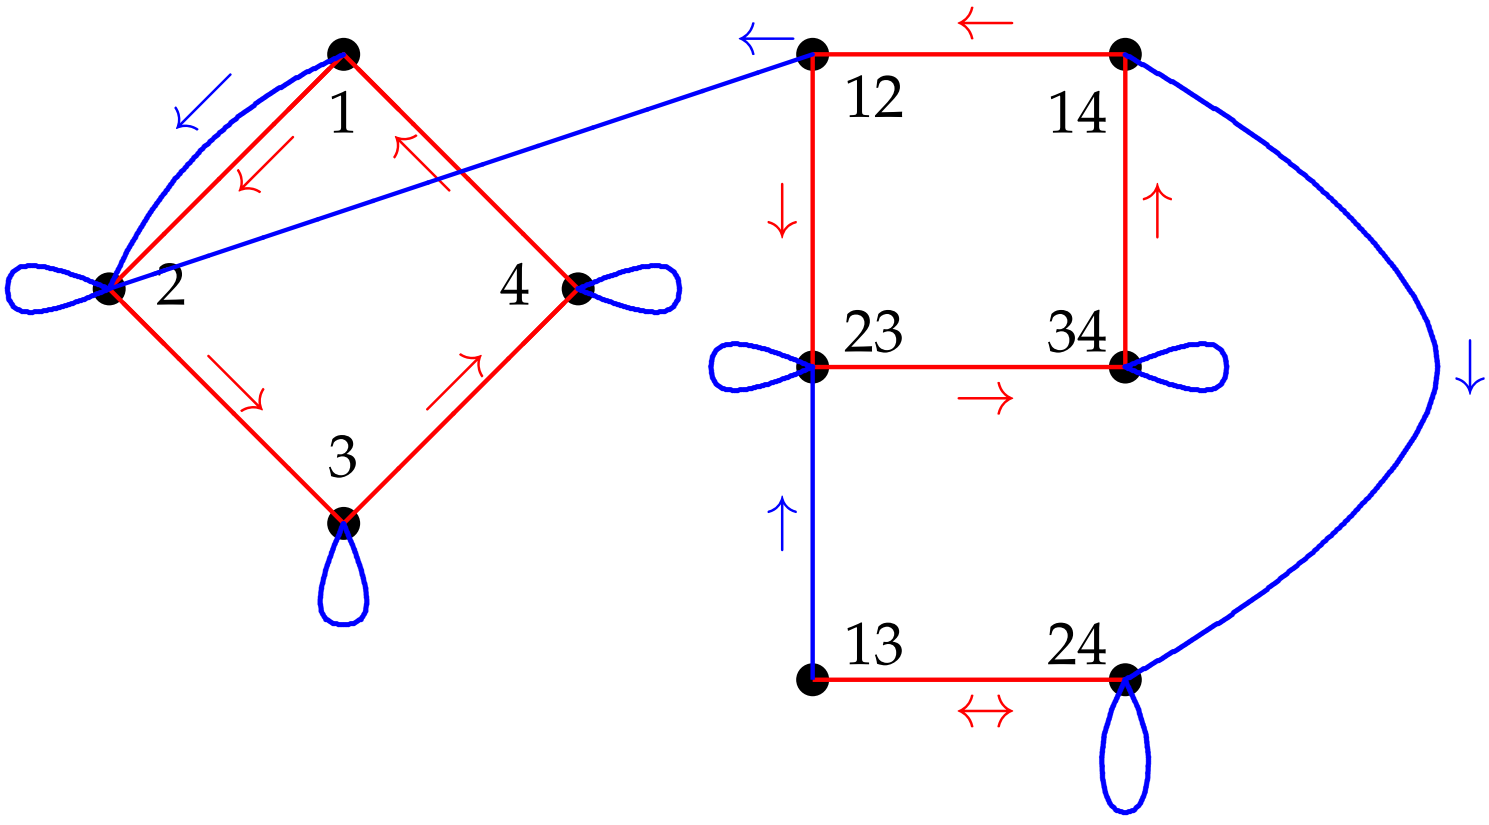

Now it suffices to check that there is a path from any vertex on the right to some vertex on the left; this can clearly be done in polynomial time.

Shortest reset word

In order to find the shortest reset word by this method, we would have to extend the diagram to all possible sets of states, and then find the shortest path from Ω to a singleton; the size of the resulting digraph would be exponentially large.
In fact:

Theorem
Deciding whether an automaton is synchronizing is in P*, but finding the length of the shortest reset word is* NP*-hard.*

The above argument gives us a cubic upper bound for the length of a reset word. For we can collapse a given pair of states in at most $\binom{n}{2}$ steps, and we only need to do this $n-1$ times to reset the automaton.

Graph homomorphisms

Let Γ and Δ be graphs. A homomorphism from Γ to Δ is a map from the vertex set of Γ to that of Δ which maps edges to edges. (We don't care what happens to non-edges; a non-edge may map to a non-edge, or to an edge, or collapse to a single vertex.)

Proposition

- *There is a homomorphism from K_r to Γ if and only if $\omega(\Gamma) \geq r$, where ω denotes clique size.*
- *There is a homomorphism from Γ to K_r if and only if $\chi(\Gamma) \leq r$, where χ denotes chromatic number.*

These things were discussed by David Roberson in his talk, but you are not expected to remember all the details.

Sandwiches

Two graphs are *homomorphically equivalent* if there are homomorphisms in both directions between them. We see from the Proposition that Γ is homomorphically equivalent to a complete graph if and only if $\omega(\Gamma) = \chi(\Gamma)$.
So for these graphs, which I will discuss much more, are those for which Simone Severini's sandwich becomes Danish ...

Graph endomorphisms

An endomorphism of a graph is a homomorphism from the graph to itself. The endomorphisms of Γ form a monoid (semigroup with identity) denoted by $\mathrm{End}(\Gamma)$.
There is a very close connection between transformation semigroups and graphs. The map $\Gamma \mapsto \mathrm{End}(\Gamma)$ takes any graph to a semigroup. We now define a map in the other direction. Given a transformation semigroup S on a set V, we define a graph $\mathrm{Gr}(S)$ on V by the rule that $v \sim w$ if and only if there is no element $f \in S$ for which $vf = wf$.
Note that, if $S \leq T$, then $\mathrm{Gr}(S)$ contains $\mathrm{Gr}(T)$ as a spanning subgraph; the map Gr is "inclusion-reversing". (The map End is not; we do not have a Galois correspondence here.)

Graphs and transformation semigroups

These correspondences have several further properties.

Theorem

Let S be a transformation semigroup on V.

- $S \leq \mathrm{End}(\mathrm{Gr}(S))$, *and* $\mathrm{Gr}(S) = \mathrm{Gr}(\mathrm{End}(\mathrm{Gr}(S)))$.
- *The three numbers $\omega(\mathrm{Gr}(S))$, $\chi(\mathrm{Gr}(S))$, and the minimal rank of an element of S are equal.*
- *S is non-synchronizing if and only if there exists a non-null graph Γ on the vertex set V with $S \leq \mathrm{End}(\Gamma)$; the graph Γ can be assumed to further satisfy $\omega(\Gamma) = \chi(\Gamma)$.*
- $\mathrm{Gr}(S)$ *is the null graph if and only if S is synchronizing, and is the complete graph if and only if $S \leq \mathrm{Sym}(V)$.*

Proof

All parts of the theorem are easy to prove; there is nothing deep here.
I will just show the last part. If G is synchronizing, then every pair can be collapsed, and the graph is null; if G consists of permutations, then no pair can be collapsed, and the graph is complete.
The only thing not quite trivial is that, if $\mathrm{Gr}(S)$ is null (that is, any pair of points can be collapsed), then S is synchronizing. Collapsing two points in the image reduces the rank, and in $n-1$ steps we reduce it to 1.

Cores and hulls

A core of a graph Γ is a graph Δ with minimum number of vertices subject to being homomorphically equivalent to Γ. It is unique up to isomorphism, and can be realised as an induced subgraph of Γ with a retraction from Γ to Δ (that is, a homomorphism which is the identity on its image).
We have seen that the graph $\mathrm{Gr}(S)$, for a transformation semigroup S, has the property that its core is complete.
A kind of dual is the hull of Γ, defined to be $\mathrm{Gr}(\mathrm{End}(\Gamma))$. It has the same vertex set as Γ.

Theorem

- Γ *is a spanning subgraph of* $\mathrm{Hull}(\Gamma)$.
- $\mathrm{End}(\Gamma) \leq \mathrm{End}(\mathrm{Hull}(\Gamma))$.
- *The core of* $\mathrm{Hull}(\Gamma)$ *is a complete graph on the vertex set of the core of* Γ.
- $\mathrm{Hull}(\mathrm{Hull}(\Gamma)) = \mathrm{Hull}(\Gamma)$.

The derived graph

We would understand synchronization much better if we could understand the maximal non-synchronizing sub-semigroups of the full transformation semigroup. In order to do that, another construction is required.
Given a graph Γ with $\omega(\Gamma) = m$, the derived graph Γ' is the graph with the same vertex set as Γ, and as edges those edges of Γ which are contained in cliques of size m.

Proposition

For any graph Γ, *we have* $\mathrm{End}(\Gamma) \leq \mathrm{End}(\Gamma')$.

The proof is straightforward.

Maximal non-synchronizing semigroups

Theorem

Let S be a maximal non-synchronizing sub-semigroup of the complete transformation semigroup on V. Then there are graphs Γ and Δ on the vertex set V satisfying

- $\mathrm{End}(\Gamma) = \mathrm{End}(\Delta) = S$;
- $\omega(\Gamma) = \chi(\Gamma) = \omega(\Delta) = \chi(\Delta)$;
- $\Gamma = \mathrm{Hull}(\Delta)$ *and* $\Delta = \Gamma'$.

I conjecture that we can take $\Gamma = \Delta$ in this theorem.

Theorem

If the non-null graph Γ on vertex set V satisfies $\Gamma = \mathrm{Hull}(\Gamma) = \Gamma'$, then $\mathrm{End}(\Gamma)$ is a maximal non-synchronizing sub-semigroup of the full transformation semigroup on V.

The probability of synchronization

Dixon's Theorem asserts that the probability that two random permutations of $\{1, \ldots, n\}$ generate the symmetric or alternating group tends to 1 as $n \to \infty$.
Two random transformations cannot generate the full transformation semigroup, but I conjectured that the probability that two random transformations generate a synchronizing semigroup tends to 1.
Mikhail Berlinkov has a paper on the arXiv (1304.5774) giving a proof of this conjecture. It uses non-trivial ideas from probability theory.
My hope is that analysis of the maximal non-synchronizing sub-semigroups can lead to another proof of this result, maybe giving more structural information.
Note that the probability that one permutation generates a transitive subgroup and the probability that one transformation generates a synchronizing semigroup are both $1/n$.

Synchronizing groups

Let G be a permutation group on V. By abuse of language, we say that G is synchronizing if the semigroup $\langle G, f \rangle$ generated by G and f is synchronizing for any non-permutation f of V.
This notion was introduced by João Araújo and Ben Steinberg; the initial hope was that it would lead to a proof of the Černý conjecture in some further cases. (Note that the examples meeting the bound are generated by a cyclic group and one further transformation.)
This has happened, but a rich theory of synchronizing groups has been developed. I will sketch some highlights.

How to recognise synchronizing groups

I will say that a subset, partition, graph, etc., on the set V is trivial if it is invariant under the symmetric group on V, and non-trivial otherwise. Thus, the trivial graphs are the complete and null graphs.

Theorem

A permutation group G on V is non-synchronizing if and only if there is a non-trivial graph Γ on V with $\omega(\Gamma) = \chi(\Gamma)$ and $G \leq \mathrm{Aut}(\Gamma)$.

This follows because if $\langle G, f \rangle \leq \mathrm{End}(\Gamma)$, then $G \leq \mathrm{Aut}(\Gamma)$.
Does this give an efficient test? See later ...

Primitive and basic groups

A permutation group G on V is said to be

- transitive if it preserves no non-trivial subset of V;
- primitive if it is transitive and preserves no non-trivial partition of V;
- basic if it is primitive and preserves no non-trivial Hamming scheme (Cartesian product structure) on V;
- 2-transitive if it preserves no non-trivial binary relation on V (and $|V| \geq 2$).

All these properties of a permutation group can be tested in polynomial time.

Synchronizing groups are basic

Theorem

- *A synchronizing group is primitive.*
- *A synchronizing group is basic.*
- *A 2-transitive group is synchronizing.*

To prove this, observe that an imprimitive group preserves a complete multipartite graph, while a non-basic group preserves a Hamming graph, and these graphs have clique number equal to chromatic number.
According to the O'Nan–Scott Theorem, a basic group is of one of three types: affine, diagonal, or almost simple.
Since we know a lot about primitive groups, following the Classification of Finite Simple Groups, it was hoped that this might lead to a classification of synchronizing groups. But things are not so easy ...

Classical groups

In many cases, testing a family of primitive basic groups for the synchronizing property leads to intractible combinatorial problems. Here is one class of examples. Apologies to those not very familiar with the terminology.

Let G be a finite classical group (symplectic, orthogonal, or unitary), acting on its associated polar space. There are just two non-trivial G invariant graphs: the orthogonality graph and its complement. It is easy to show, considering cliques and colourings of these graphs, that G is non-synchronizing if and only if *either*

- the polar space contains an ovoid and a spread of maximal subspaces; *or*
- the polar space has a partition into ovoids.

Despite decades of work by finite geometers, we still do not have a complete list of which classical polar spaces have either of the structures described. Can the new techniques for deciding if strongly regular graphs are cores resolve this?

Testing synchronization

As we saw, the property of being synchronizing is stronger than those of being transitive, primitive or basic, and weaker than 2-transitivity. In contrast to these, we do not have an efficient test. The best we can do to test a given group G is

- construct all non-trivial G-invariant graphs (there are $2^r - 2$ of these, where r is the number of G-orbits on 2-sets);
- test each of these graphs Γ to find whether its clique number and chromatic number are equal.

This involves exponentially many NP-hard problems! But the graphs are rather special, since they admit basic primitive groups. Permutation groups with degrees in the thousands have been tested by this method.

Synchronizing non-uniform maps

I conclude with two important open problems. In each case the problem suggests that, while there are primitive non-synchronizing groups, primitivity is in some sense the important watershed.
The first is due to João Araújo. A transformation of V is called non-uniform if not all inverse images of points in its image have the same cardinality. Also, we say that the permutation G synchronizes the map f if $\langle G, f \rangle$ is synchronizing.

Conjecture

Let G be a primitive permutation group. Then G synchronizes every non-uniform map.

The first case is a theorem of Rystsov, which asserts that a permutation group of degree n is primitive if and only if it synchronizes every map of rank $n-1$. This is easily deduced by graph-theoretic methods.

Non-synchronizing ranks

Let G be a permutation group of degree n. A number r with $2 \leq r \leq n-1$ is an non-synchronizing rank of G if there is a map of rank r not synchronized by G.

Proposition

An imprimitive group of degree n has at least $(\frac{3}{4} - o(1))n$ non-synchronizing ranks.

Conjecture

A primitive group of degree n has at most $O(\log n)$ non-synchronizing ranks, while a basic group has at most $O(\log \log n)$ non-synchronizing ranks.

The automorphism group of the Hamming scheme $H(r, m)$, namely $S_m \wr S_r$, has non-synchronizing ranks m^i for $i = 1, 2, \ldots, r-1$, and probably no others (work of my student Artur Schaefer). Indeed, for the case $r = 2$, Godsil and Royle showed that the only non-synchronizing rank is m.

Thank you – and happy birthday Chris!