Endomorphisms and synchronization, 3: The almost synchronizing conjecture

Peter J. Cameron

BIRS, November 2014

Recapitulation

A permutation group G synchronizes a map f if $\langle G, f \rangle$ is synchronizing (i.e. contains a constant function).
A permutation group G is synchronizing if it synchronizes every non-permutation, and is almost synchronizing if it synchronizes every non-uniform map.
Rystsov's Theorem says that a permutation group of degree n is primitive if and only if it synchronizes every map of rank $n - 1$. Such a map is necessarily non-uniform (if $n > 2$). So, for $n > 2$, we have

$$\text{synchronizing} \Rightarrow \text{almost synchronizing} \Rightarrow \text{primitive}.$$

Further implications

Fof $n > 2$, we have

$$\text{primitive} \Rightarrow \text{transitive}.$$

A permutation group G on X is 2-transitive if it acts transitively on ordered pairs of distinct elements of X, and is 2-homogeneous if it acts transitively on 2-element subsets of X. For $n > 2$, we have

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

The second implication holds because, as we saw, a monoid M is synchronizing if any pair of points can be identified by any point of M; if f is not a permutation, it identifies some pair x and y, and if G is 2-homogeneous then any pair can be mapped to $\{x, y\}$ and then collapsed by f.

Applications of CFSG

The classification of finite simple groups (CFSG) has resulted in a great increase in our understanding of permutation groups. In particular, all 2-transitive groups are explicitly known. Also, it had been shown earlier that 2-homogeneous groups are 2-transitive, with explicitly known exceptions.
We understand primitive groups a lot better, but are still a long way from a complete classification. However, all primitive groups up to degree 4095 have been determined, and are available in the computer systems Magma and (hopefully soon!) GAP.
Synchronizing groups, however, are more mysterious.
We saw in the last lecture Araújo's conjecture that "almost synchronizing" is equivalent to "primitive".

Pseudocores

A pseudocore is a graph Γ with core K_r (that is, $\omega(\Gamma) = \chi(\Gamma) = r$) with the property that every endomorphism which is not an automorphism has image a core of Γ (and so is an r-colouring of Γ).

If G is a primitive group and every G-invariant graph is either a core or a pseudocore, then Araújo's conjecture holds for G.

There are quite a lot of examples, including many of the "large" permutation groups.

It often happens that endomorphisms of a pseudocore correspond to important and prolific combinatorial structures.

Example

- The line graph of K_m is a core if m is odd, a pseudocore if m is even; its complement is a core (an easy consequence of the theorem of Lovász on the chromatic number of Kneser graphs, though it can be proved directly.)
  Proper endomorphisms of $L(K_m)$ for m even come from 1-factorisations of K_m with $m-1$ colours. There are more than exponentially many of these (in terms of $n = m(m-1)/2$).
- The $m \times m$ grid graph (the line graph of $K_{m,m}$) and its complement are both pseudocores.
  Proper endomorphisms of $L(K_m, m)$ come from Latin squares of order m. There are more than exponentially many of these (in terms of $n = m^2$).

So the "large" primitive groups S_m on 2-sets and $S_m \wr S_2$ both satisfy Araújo's conjecture.

Large rank

João Araújo, Wolfram Bentz and I began attacking the conjecture, and showed the following two results.

Proposition

A primitive group of degree n synchronizes every map of rank $n-4$ or greater.

The proof follows the lines we saw in Rystsov's theorem but is much more complicated:

- For each rank $n-k$, the number of possible kernel types of a map depends only on k (for large n; for small n some may fail to exist). For example, for rank $n-3$ the possible kernel types are $4,1,1,\ldots,1$; $3,2,1,\ldots,1$; or $2,2,2,1,\ldots,1$.
- For each kernel type, carefully examine the possible configurations of edges between the non-trivial kernel classes and how they map under f.
- Use the fact that if a primitive group of degree n (other than S_n or A_n) contains an element moving m points, then n is bounded above by a function of m.

Small rank

Proposition

A primitive group synchronizes every non-uniform map of rank at most 4.

I will sketch some ideas in the proof since they are relevant to what follows.

First, an easy result of Peter Neumann. Suppose G fails to synchronize a map of rank 2. Then G is a group of automorphisms of a graph Γ with chromatic number 2 (i.e. a bipartite graph). If Γ is disconnected, its connected components are blocks of imprimitivity; while a connected bipartite graph has a unique bipartition, so again its automorphism group is imprimitive.

Primitive graphs

What properties distinguish a primitive graph (one with primitive automorphism group)?
There cannot be a purely local property, since every finite graph occurs as an induced subgraph of a primitive graph.

Proposition

A primitive graph with chromatic number r cannot have a subgraph isomorphic to K_{r+1} with an edge removed.

Proof.

Suppose the contrary, and let $\{v, w\}$ be the removed edge. It is clear that v and w have the same colour in any r-colouring of the graph. Thus, "have the same colour in any r-colouring" is a non-trivial equivalence relation invariant under the automorphism group of the graph, a contradiction to primitivity. $\square$

Consequences

Suppose that G is primitive but not synchronizing; let Γ be a G-invariant graph with clique number and chromatic number r, and let f be an endomorphism of rank r. We have seen earlier that f is uniform.

Proposition

Γ *has no endomorphism of rank* $r+1$.

Proof.

If g were such an endomorphism, then (up to permutation) one kernel class of f would split into two classes for g, and the other classes would remain the same. So the induced subgraph on the image of g would be K_{r+1} with an edge removed. □

Corollary

A primitive group synchronizes every non-uniform map of rank at most 4.

For in a counterexample, the minimum rank of an endomorphism would be 2 or 3, both of which are impossible.

Butterfly or bowtie?

A significant role in what comes next is played by the following graph:

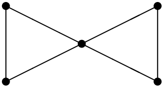

You may know this graph as the bowtie; I will call it the butterfly, since you will see that by a flap of its wings it introduces chaos into the theory of synchronization!

Rank 5

If a primitive graph Γ has a non-uniform endomorphism of rank 5, then it has clique and chromatic number 3, and the image of the rank 5 endomorphism does not contain $K_4 - e$. The obvious candidate is the butterfly:

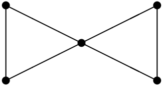

So we require this to be a subgraph of Γ. The most economical way to do this is to require it to be the closed neighbourhood of a vertex.

The study of primitive groups in which the stabiliser of a point has an orbit of size 4 was begun by Charles Sims in the late 1960s, and was completed by Cai Heng Li, Zai Ping Lu and Dragan Marušič in 2004. There are just three primitive graphs which have vertex neighbourhoods of the type we require:

- the flag graph of the projective plane of order 2 (the Fano plane), aka the line graph of the Heawood graph, with 21 vertices;
- the flag graph of the generalised quadrangle of order 2, aka the line graph of the Tutte–Coxeter graph, with 45 vertices;
- the line graph of the Biggs–Smith graph, with 153 vertices.

All these graphs are line graphs of trivalent graphs of high girth, and so have clique number 3.

We are looking for a graph with chromatic number 3. We know that a 3-colouring must be uniform. So we can proceed as follows:

- find an independent set of size $n/3$;
- test whether the induced subgraph Δ on its complement is bipartite.

In each of our three examples, this process succeeds. Indeed, Δ has valency 2 in each case.
For the first graph, there is only one independent set of size 7 (up to symmetry), and its complement is a 14-cycle. So the colouring does not factor through an endomorphism of larger rank.

A counterexample!

For the second graph, there is an independent set A of size 15 whose complement consists of a 10-cycle and a 20-cycle. If $\{B, C\}$ is the bipartition of the 10-cycle and $\{D, E\}$ of the 20-cycle, then we have a homomorphism of Γ onto the butterfly, mapping A to the body, B and C to the vertices on one wing, and D and E to the other.

So the primitive group $\mathrm{Aut}(\Gamma)$ is not almost synchronizing; it fails to synchronize a map with kernel classes of sizes $5, 5, 10, 10, 15$.

Further analysis showed that the 45-graph also has endomorphisms of rank 7 whose image is the “double butterfly”:

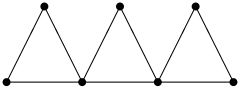

The kernel classes have sizes $10, 10, 5, 5, 5, 5, 5$. Endomorphisms of rank 5 arise by folding in one wing, and of rank 3 by folding in both wings.

The constraint satisfaction software MINION counts the endomorphisms very quickly. As well as the 1440 automorphisms, there are 25920 endomorphisms of rank 7, 51840 of rank 5, and 25920 of rank 3.

The endomorphism monoid of the graph is generated by the automorphism group together with one endomorphism of rank 7.

Note on computation

It is well known that Graph Automorphism is a candidate for a problem of intermediate complexity between P and NP-complete: it lies in NP but has not been shown to lie in either of the above classes.
By contrast, Graph Endomorphism is NP-complete.
The way the above computation was first performed was as follows:

- First use MINION to count the endomorphisms.
- Then use GAP to find them. Each time we find one (by extending a map one point at a time), compute the monoid generated by what we already found; the next candidate can be outside this monoid.
- Stop when we have found them all (so exhaustive backtrack not required).

So a combination of algebraic and constraint-satisfaction software is good …

The 153-graph

The 153-vertex primitive graph Γ, the line graph of the Biggs–Smith graph, is even more interesting.
We have not been able to find all the endomorphisms of Γ. But Gordon Royle has discovered many strange endomorphisms.
There are endomorphisms of rank 9, with kernel classes $39, 39, 39, 6, 6, 6, 6, 6, 6$. Its image is a triangle with triangles attached at each vertex:

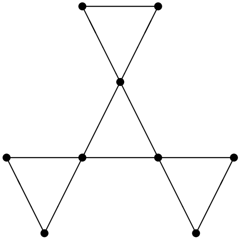

There are endomorphisms of rank 7 whose image is the double butterfly, with kernel classes

- $45, 33, 27, 18, 18, 6, 6$
- $45, 34, 28, 17, 17, 6, 6$
- $45, 36, 30, 15, 15, 6, 6$
- $45, 37, 31, 14, 14, 6, 6$
- $45, 38, 32, 13, 13, 6, 6$
- $45, 39, 33, 12, 12, 6, 6$
- $45, 40, 34, 11, 11, 6, 6$
- $45, 42, 36, 9, 9, 6, 6$
- $45, 43, 37, 8, 8, 6, 6$
- $45, 45, 39, 6, 6, 6, 6$

Only the last of these can arise by folding in one triangle of the image of the rank 9 map. So it appears that we cannot generate the endomorphism monoid with automorphisms and one more element.

For rank 5, the following kernel types occur:

$51, 26, 26, 25, 25$	$51, 36, 36, 15, 15$
$51, 27, 27, 24, 24$	$51, 37, 37, 14, 14$
$51, 28, 28, 23, 23$	$51, 38, 38, 13, 13$
$51, 29, 29, 22, 22$	$51, 39, 39, 12, 12$
$51, 30, 30, 21, 21$	$51, 40, 40, 11, 11$
$51, 31, 31, 20, 20$	$51, 41, 41, 10, 10$
$51, 32, 32, 19, 19$	$51, 42, 42, 9, 9$
$51, 33, 33, 18, 18$	$51, 43, 43, 8, 8$
$51, 34, 34, 17, 17$	$51, 45, 45, 6, 6$
$51, 35, 35, 16, 16$	

And, of course, endomorphisms of rank 3 with kernel classes $51, 51, 51$ whose image is a triangle.

Further examples

We can relax the restriction that the graph has valency 4 (and the closed neighbourhood is a butterfly) – we only require that it contains a butterfly.

There are two graphs of valency 6 on 495 vertices, each with automorphism group $M_{12} : 2$, in which a closed neighbourhood consists of three triangles with a common vertex. These graphs are too large to search for independent sets of size 165; but we can cheat.

In each case, there is a subgroup $\mathrm{PSL}(2, 11)$ of the automorphism group with orbits of sizes $55, 55, 110, 110, 165$. It can be verified that each orbit is an independent set, and the "collapsed graph" is the butterfly. Since each graph contains a butterfly, there exist non-uniform endomorphisms of rank 5.

Problems

Problem

- *Do there exist infinitely many primitive groups which fail to synchronize maps of rank* 5*?*
- *We've seen that if the minimum rank of an endomorphism of a primitive graph is r, there cannot be an endomorphism of rank $r + 1$. Can a primitive graph have endomorphisms with ranks differing by* 1*? Can a primitive group fail to synchronize maps with ranks differing by* 1*?*
- *What about higher rank? (see next slide ...)*

Rank 6

We have two infinite families of primitive graphs with the property that the minimum rank of an endomorphism is 4 and there are non-uniform endomorphisms of rank 6 with kernel type $2^{m-2}, 2^{m-2}, 2^{m-3}, 2^{m-3}, 2^{m-3}, 2^{m-3}$. The image is a larger "butterfly" consisting of two K_4s with a common edge.

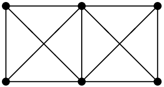

The graphs are Cayley graphs for elementary abelian 2-groups of order 2^m for various m.
The smallest such graph has 64 vertices.
One of these graphs, on 256 vertices, also has non-uniform endomorphisms of ranks 9 (with kernel type $64, 32, 32, 32, 16, 16, 16, 16$) and 12 (with kernel type 32 (4 times) and 16 (8 times).

Large rank revisited

Since the conjecture about almost synchronizing groups is false, it is perhaps worth making a weaker conjecture to guide investigations. We propose the following.

Conjecture

A primitive group of rank n synchronizes any map of rank r with $n/2 < r < n$.

Note that such a map is automatically non-uniform.

Artur Schaefer has shown that a primitive group with degree n and permutation rank 3 (that is, three orbits on ordered pairs) synchronizes a non-permutation of rank r if $r > n - c\sqrt{n}$ (where we can take $c = 1/24$). The proof uses the fact that all such groups are known (following CFSG).

Non-synchronizing ranks

Part of the rationale behind Araújo's conjecture was that, even though the properties "primitive" and "synchronizing" are not equivalent, there is a big difference between primitive and imprimitive groups with respect to synchronization. Here is one way to quantify this.

Let G be a permutation group. Define the set $\mathrm{NS}(G)$ of non-synchronizing ranks of G to be the set of all $r < n$ for which there exists a map of rank r which is not synchronized by G.

The meta-conjecture is that imprimitive groups have many non-synchronizing ranks but primitive groups have relatively few.

Imprimitive groups

Let G be transitive but imprimitive of degree n. Suppose that some G-invariant equivalence relation has m blocks of size k, with $m, k > 1$ and $mk = n$.
Then G preserves a disjoint union of m complete graphs of size k, and its complement, the complete m-partite graph with parts of size k.
The first of these has endomorphisms of any rank which is a multiple of k. The second has endomorphisms of any rank r satisfying $m < r < n$, since we may collapse arbitrarily in each part.
Together these show that $|\mathrm{NS}(G)| \geq (\frac{3}{4} - o(1))n$.

Primitive groups

Conjecture

If G is primitive of degree n, then $|\mathrm{NS}(G)| = O(\log n)$.

Non-basic primitive groups of degree $n = m^k$ have non-synchronizing ranks n^i for $1 \leq i \leq k-1$. So the conjecture, if true, is best possible.

However, it may be that a stronger result holds for basic primitive groups.

Do any of these newly discovered examples provide a disproof of this conjecture?

The infinite

I have mostly avoided talking about the infinite, since the other tutorials cover this. Here is an example, showing that at least some aspects of the infinite are easier than the finite.

Let Γ be the line graph of the infinite trivalent tree. Then Γ has clique number and chromatic number 3 (indeed, it has $2^{\aleph_0}$ 3-colourings), and its automorphism group is (strongly) primitive.

Choose a fixed 3-colouring of Γ, and an induced subgraph Δ which is connected and triangle-closed. We can retract Γ onto Δ as follows:

- interior vertices of Δ are fixed;
- vertices in the component attached at a boundary triangle T are mapped onto T, using the given 3-colouring.

Thus Γ has endomorphisms of every finite odd rank.

Separation

A permutation group fails to be synchronizing if there is a non-trivial subset A and partition P of the domain with the property that, for every part B of P and every $g \in G$, we have $|Ag \cap B| = 1$. (The map that takes a point $x \in B$ to the unique point of $A \cap B$ for every $B \in P$ is not synchronized by P.)

We say that a permutation group is non-separating if there are non-trivial subsets A and B such that, for all $g \in G$, $|Ag \cap B| = 1$; it is separating otherwise.

A transitive permutation group G is non-separating if and only if there is a non-trivial G-invariant graph Γ with $|\omega(\Gamma)| \cdot |\omega(\overline{\Gamma})| = n$.

Thus, "separating" implies "synchronizing". The converse is false but only a few counterexamples are known: mainly automorphism groups of polar spaces which have an ovoid but no spread and no partition into ovoids.

More?

There is plenty more to say about synchronization. We have a couple of related properties intermediate between "separating" and "2-homogeneous". There are also connections with representation theory over subfields of the complex numbers. Many of the problems I discussed in lectures on synchronization in 2010 are still open: see

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

The penultinmate chapter of the notes concerns the infinite, and the last chapter has open problems.
But that is enough for now … Thank you for your attention!