EMS Surv. Math. Sci. x (201x), xxx–xxx
DOI 10.4171/EMSS/x

EMS Surveys in Mathematical Sciences

Between primitive and 2-transitive: Synchronization and its friends

João Araújo * Peter J. Cameron ** Benjamin Steinberg ‡

Contents

*The first author was supported by the Fundação para a Ciência e Tecnologia (Portuguese Foundation for Science and Technology) through the project CEMAT-CIÊNCIAS UID/Multi/04621/2013, and through project Hilbert's 24th problem PTDC/MHC-FIL/2583/2014.

**The second author was supported by the Fundação para a Ciência e Tecnologia (Portuguese Foundation for Science and Technology) through the project CEMAT-CIÊNCIAS UID/Multi/04621/2013.

‡The third author was supported in part by a grant from the Simons Foundation (#245268 to Benjamin Steinberg), the Binational Science Foundation of Israel and the US (#2012080), a CUNY Collaborative Research Incentive Grant #2123 and by a PSC-CUNY grant.

J. Araújo, Universidade Aberta and CEMAT-CIÊNCIAS, Universidade de Lisboa, Campo Grande, C6, 1749-016 Lisboa, Portugal
E-mail: jjaraujo@fc.ul.pt

P. J. Cameron, School of Mathematics and Statistics, University of St Andrews, St Andrews, Fife KY16 9SS, UK, and CEMAT-CIÊNCIAS, Universidade de Lisboa, Portugal
E-mail: pjc20@st-andrews.ac.uk

B. Steinberg, Department of Mathematics, City College of New York, New York, NY 10031, USA
E-mail: bsteinberg@ccny.cuny.edu

An automaton (consisting of a finite set of states with given transitions) is said to be synchronizing if there is a word in the transitions which sends all states of the automaton to a single state. Research on this topic has been driven by the *Černý conjecture*, one of the oldest and most famous problems in automata theory, according to which a synchronizing n-state automaton has a reset word of length at most $(n-1)^2$. The transitions of an automaton generate a transformation monoid on the set of states, and so an automaton can be regarded as a transformation monoid with a prescribed set of generators. In this setting, an automaton is synchronizing if the transitions generate a constant map.

A permutation group G on a set Ω is said to synchronize a map f if the monoid $\langle G, f\rangle$ generated by G and f is synchronizing in the above sense; we say G is synchronizing if it synchronizes every non-permutation.

The classes of synchronizing groups and friends form an hierarchy of natural and elegant classes of groups lying strictly between the classes of primitive and 2-homogeneous groups. These classes have been floating around for some years and it is now time to provide a unified reference on them. The study of all these classes has been prompted by the Černý conjecture, but it is of independent interest since it involves a rich mix of group theory, combinatorics, graph endomorphisms, semigroup theory, finite geometry, and representation theory, and has interesting computational aspects as well. So as to make the paper self-contained, we have provided background material on these topics.

Our purpose here is to present recent work on synchronizing groups and related topics. In addition to the results that show the connections between the various areas of mathematics mentioned above, we include a new result on the Černý conjecture (a strengthening of a theorem of Rystsov), some challenges to finite geometers (which classical polar spaces can be partitioned into ovoids?), some thoughts about infinite analogues, and a long list of open problems to stimulate further work.

Mathematics Subject Classification (2010). 20B15; 20M20

Keywords. Permutation groups, transformation semigroups, automata, synchronization, primitivity.

Acknowledgment. We are grateful to the referee, whose thoughtful and careful report has substantially improved this paper.

1. Introduction

The study of primitive and multiply-transitive permutation groups is one of the oldest parts of group theory, going back to Jordan and Mathieu in the nineteenth century.

Recently, in the study of synchronization of automata, various other classes of permutation groups have been considered, most notably the *synchronizing groups*, those permutation groups which, together with any transformation which is not a permutation, generate a constant map. The class of synchronizing groups contains the 2-transitive (or 2-homogeneous) groups, and is contained in the class of primitive groups (or indeed the class of *basic* groups in the O'Nan–Scott classification).

The purpose of this paper is to introduce the theories of synchronizing groups, and of various related classes of groups. We stress the links to semigroup theory and automata theory and give numerous examples to show that for the most part all our classes of groups are distinct.

Since the paper covers a wide variety of subject matter, we have taken some trouble to include enough background material to make it self-contained.

After a brief introduction we give in Section 1 some theory of permutation groups, transformation monoids, graphs and digraphs. Section 2 further develops the theory of permutation groups, introducing the notions of transitive, primitive, 2-homogeneous and 2-transitive groups, the O'Nan–Scott theorem, and the Classification of Finite Simple Groups.

In Section 3 we introduce our main concern, the notion of synchronization for finite automata, which can be expressed as a property of transformation monoids. We introduce one of the main research problems in this area, the famous Černý conjecture. We define the class of synchronizing groups, and study its relation to the classes already defined.

Section 4 introduces graph homomorphisms and endomorphisms, and characterizes synchronizing monoids and groups in terms of these concepts. Section 5 defines several related classes of permutation groups. Section 6 gives some examples, concentrating on the action of the symmetric group on k-sets and the action of a classical group on its polar space.

Section 7 links some of the properties of permutation groups we have considered with representation theory, and introduces some new classes. Section 8 gives alternative characterizations of some of our classes in terms of functions. Section 9 explains the connection between some of our results and instances of the Černý conjecture, including a new theorem which strengthens a result of Rystsov.

In Section 10, we look at other classes of permutation groups lying between primitive and 2-transitive, and in Section 11 we take a brief look at what happens in the infinite case. The final Section 12 lists a number of unsolved problems.

In the remainder of this section, we give a brief outline to permutation groups, transformation monoids, graphs and digraphs.

This survey grew from a course given by the second author in 2010 at the London Taught Course Centre; we are grateful to the course participants for their comments.

Apart from the present authors, many others have contributed to the theory presented here. We are particularly grateful to Peter Neumann, whose contributions are discussed in many places. Csaba Schneider, Leonard Soicher and Pablo Spiga wrote GAP code for determining which primitive permutation groups of small degree are synchronizing, which guided many of our conjectures. Others to whom we are grateful include Wolfram Bentz, Michael Brough, Ian Gent, Nick Gravin, Ferdinand Ihringer, Cristy Kazanidis, Tom Kelsey, James Mitchell, Dima Pasechnik, Colva Roney-Dougal, Gordon Royle, Nik Ruškuć, Jan Saxl and Artur Schaefer.

1.1. Permutation groups. For general references on permutation groups, we recommend [35, 47, 113].

The *symmetric group* $\mathrm{Sym}(\Omega)$ on a set Ω is the group whose elements are all the permutations of Ω and whose operation is composition. If $\Omega = \{1, 2, \ldots, n\}$, we write the group as S_n, the symmetric group of degree n. We write permutations on the right of their argument, so that ag is the image of a under the permutation g: this has the advantage that the composition "g followed by h" is gh, rather than hg.

A *permutation group* G on Ω is a subgroup of $\mathrm{Sym}(\Omega)$. The *degree* of G is the cardinality of Ω.

Almost always in this paper, Ω is a finite set.

It is more usual to define an *action* of a group G on a set Ω, this being defined as a homomorphism from G to the symmetric group $\mathrm{Sym}(\Omega)$. The advantage is that the same group may act on several different sets. Note that the image of an action is a permutation group. Most concepts defined below, starting with transitivity, can be extended to group actions by saying that (for example) an action is transitive if its image is a transitive permutation group.

It is well known from elementary discrete mathematics that a permutation on a finite set can be decomposed into disjoint cycles. A similar decomposition applies to a permutation group. Let G be a permutation group on Ω. Define an equivalence relation $\equiv$ on Ω by the rule that $a \equiv b$ if $ag = b$ for some $g \in G$. (The reflexive, symmetric and transitive laws for $\equiv$ follow immediately from the identity, inverse and closure axioms for G.) The equivalence classes are the *orbits* of G on Ω.

We say that G is *transitive* if it has just one orbit.

1.2. Transformation monoids. The set of all mappings from Ω to itself, with the operation of composition, is a *monoid*: that is, it is closed and associative and has an identity element. It is called the *full transformation monoid* on Ω, denoted by $T(\Omega)$, or T_n if $\Omega = \{1, \ldots, n\}$. As for permutations, we write a transformation on the right of its argument.

Note that $\mathrm{Sym}(\Omega)$ is a subgroup of $T(\Omega)$. The difference $T(\Omega) \setminus \mathrm{Sym}(\Omega)$ consists of all the *singular maps* on Ω.

Let $f \in T(\Omega)$. The *image* of f, which we write as $\mathrm{Im}(f)$ or Ωf, is the subset

$\{af : a \in \Omega\}$ of Ω. The *rank* of f is the cardinality of its image. The *kernel* $\mathrm{Ker}(f)$ of f is the equivalence relation $\equiv$ on Ω defined by $a \equiv b$ if $af = bf$; we will not distinguish between the equivalence relation and the corresponding partition. Note that the number of equivalence classes of $\mathrm{Ker}(f)$ is equal to the rank of f.

1.3. Graphs and digraphs. A *graph* on the vertex set Ω can be regarded in several ways: as a symmetric binary relation called *adjacency* on Ω, or as a collection of subsets called *edges*, each of cardinality 1 or 2. Note that this definition forbids multiple edges. Usually we will also forbid loops: that is, we do not allow a vertex to be adjacent to itself (so edges cannot have cardinality 1).

Let Γ be a graph on the vertex set Ω. An *induced subgraph* of Γ is obtained by choosing a subset A of Ω as vertex set, and including all edges of Γ which join two vertices of A. A *spanning subgraph* is obtained by choosing the whole of Ω as vertex set, but taking a subset of the edge set of Γ. Thus, for example, the Petersen graph (Figure 1) does not contain a 4-cycle as induced subgraph, but does have the disjoint union of two 5-cycles as a spanning subgraph.

A graph is *connected* if, given any two vertices v and w, there is a sequence $v = x_0, x_1, \ldots, x_d = w$ of vertices such that x_{i-1} and x_i are adjacent for $i = 1, \ldots, d$. The smallest such d (for given v and w) is the *distance* from v to w.

The *complete graph* on a given vertex set has all possible edges; the *null graph* has no edges. These graphs are denoted by K_n and N_n if there are n vertices. The *line graph* of a graph Γ is the graph $L(\Gamma)$ whose vertex set is the edge set of Γ, two vertices of $L(\Gamma)$ being adjacent if the corresponding edges of Γ have a common vertex. The *complement* $\overline{\Gamma}$ of Γ is the graph with the same vertex set as Γ, two vertices being adjacent in the complement if and only if they are not adjacent in Γ.

Let P be a partition of a set Ω. The *complete multipartite graph* on Ω with multipartition P is the graph in which two vertices are adjacent if and only if they belong to different parts of P. If the number of parts is 2, we speak of a *complete bipartite graph* with bipartition P. A spanning subgraph of a complete multipartite graph is called *multipartite*, and similarly for bipartite.

An *automorphism* of a graph is a permutation of the vertex set which maps edges to edges. (If the graph is infinite, we must also require that it maps non-edges to non-edges.) The set of all automorphisms of a graph is a group, a permutation group on the vertex set, called the *automorphism group* of the graph.

The best known graph is the *Petersen graph*, the graph with ten vertices and fifteen edges shown in Figure 1.

The Petersen graph has 120 automorphisms, forming a group isomorphic to the symmetric group S_5. (To prove this, one first argues directly that there cannot be more than 120 automorphisms. Then label the vertices by pairs of elements from $\{1, \ldots, 5\}$ in such a way that two vertices are adjacent if and only if their labels are disjoint. In other words, the Petersen graph is the complement of the line graph of K_5. So the symmetric group on $\{1, \ldots, 5\}$ has an action on the vertices which preserves the adjacency relation.)

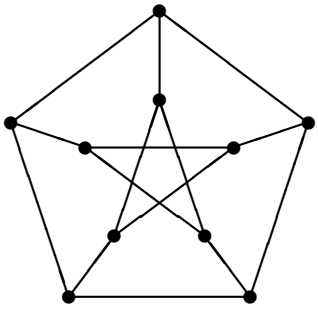

Figure 1. The Petersen graph

A *directed graph*, or *digraph*, is similarly defined, except that an edge is an ordered pair of vertices. So each edge has a direction, say from v to w, and can be represented by an arrow with tail at v and head at w. Such directed edges are called *arcs*.

We say that a directed graph is *connected*, if, when we ignore the directions of the arcs (and keep only one of any pair of edges thus created), the undirected graph so obtained is connected. A directed graph is *strongly connected* if, given any two vertices v and w, there is a sequence $v = x_0, x_1, \ldots, x_d = w$ of vertices such that there is an arc from x_{i-1} to x_i for $i = 1, \ldots, d$.

Automorphisms of directed graphs are defined similarly to the undirected case. The following important result holds.

Theorem 1.1. *Let D be a finite directed graph whose automorphism group is transitive on the vertices. If D is connected, then it is strongly connected.*

Proof. Let $R(x)$ be the set of vertices which can be reached by directed paths starting at x. Clearly we have

(a) $x \in R(x)$;

(b) if $y \in R(x)$, then $R(y) \subseteq R(x)$;

(c) if an automorphism g carries x to y, then it maps $R(x)$ to $R(y)$.

Now by the third property, if $\mathrm{Aut}(D)$ is vertex-transitive, then $|R(x)|$ is constant for all vertices x; by the second property, if $y \in R(x)$, then $R(y) = R(x)$, and so by the first property $x \in R(y)$.

Now suppose that D is connected. Given v and w, take a path from v to w in the undirected graph. Now by what we have proved, any arc in the "wrong" direction can be replaced by a path all of whose arcs are in the "right" direction; so D is strongly connected. □

2. Transitivity and primitivity

In this section, we introduce some of the basic notions of permutation group theory, especially primitivity.

We begin this section with a convention that will prove useful later. We say that a structure on a set Ω is *trivial* if it is preserved by the whole symmetric group on Ω, and is *non-trivial* otherwise.

So, for example,

(a) the trivial subsets of Ω are the empty set and the whole of Ω;

(b) the trivial partitions of Ω are the partition all of whose parts are singletons (corresponding to the equivalence relation of equality), and the partition with a single part;

(c) the trivial graphs on Ω are the complete and null graphs.

2.1. Transitivity. As we saw in the last section, a permutation group G on Ω is *transitive* if we can map any element of Ω to any other by an element of G. In the convention introduced above, G is transitive if the only G-invariant subsets of Ω are the trivial ones. The same definition applies to an action of G on Ω.

The *stabilizer* G_a of a point $a \in \Omega$ is the set

$$\{g \in G : ag = a\}$$

of elements of G (which is easily seen to be a subgroup of G).

There is an internal description of the transitive actions of a group, as follows. Let H be a subgroup of G. The *coset space* $H\backslash G$ consists of all right cosets Hx of H in G; there is an action of G on $H\backslash G$, where the permutation corresponding to the group element g maps the coset Hx to the coset Hxg. The action of G on $H\backslash G$ is transitive, and the stabilizer of the coset H is the subgroup H.

Two actions of G on sets Ω_1 and Ω_2 are *isomorphic* if there is a bijection $\phi\colon \Omega_1 \to \Omega_2$ commuting with the action of G, that is, such that $(ag)\phi = (a\phi)g$ for all $a \in \Omega_1$ and $g \in G$.

Theorem 2.1. (a) *Any transitive action of G is isomorphic to the action of G on a coset space (specifically, on $H\backslash G$, where H is the stabilizer of a point).*

(b) *The actions of G on coset spaces $H\backslash G$ and $K\backslash G$ are isomorphic if and only if H and K are conjugate subgroups of G.*

In particular, a group G has a unique (up to isomorphism) *regular* action, characterized by the fact that it is transitive and the stabilizer of a point is the identity. If we identify a singleton subset of G (a coset of the identity subgroup) with an element, this is the action of G on itself by right multiplication, as used by Cayley to show that every group is isomorphic to a permutation group.

2.2. Primitivity. A transitive permutation group G on Ω is said to be *primitive* if the only G-invariant partitions of Ω are the trivial ones.

Primitivity is a very important concept in permutation group theory, and we will see several further characterizations of it.

A subset B of Ω is called a *block*, or *block of imprimitivity*, for G if, for all $g \in G$, either $Bg = B$ or $Bg \cap B = \emptyset$.

Proposition 2.2. *The transitive permutation group G on Ω is primitive if and only if the only blocks for G are the empty set, singletons, and the whole of Ω.*

Proof. A part of any G-invariant partition is clearly a block. Conversely, if B is a non-empty block, then for all $g, h \in G$, we have $Bg = Bh$ or $Bg \cap Bh = \emptyset$; so the translates of B under G form a G-invariant partition. The result follows. □

We saw in the last subsection that any transitive group G can be identified with G acting on a coset space $H\backslash G$.

Proposition 2.3. *The action of G in $H\backslash G$ is primitive if and only if H is a maximal subgroup of G.*

Proof. If $H \leq K \leq G$, then the cosets of H contained in K form a block for G; every block containing the coset H arises in this way. □

2.2.1. Normal subgroups. As Cayley observed, every group is isomorphic to a transitive permutation group. However, not every group is isomorphic to a primitive permutation group; primitive groups have strong restrictions on their normal subgroup structure. The basic observation is:

Proposition 2.4. *A non-trivial normal subgroup of a primitive group is transitive.*

Proof. This follows from the observation that the orbits of a normal subgroup of a transitive group are blocks for the group. □

Theorem 2.5. *A primitive permutation group has at most two minimal normal subgroups; if there are two, then they are isomorphic, non-abelian, and regular, and each is the centralizer of the other in the symmetric group.*

Proof. A permutation group (not necessarily transitive) is called *semiregular* if the stabilizer of any point is the identity. Thus a transitive semiregular group is regular. It is easy to show that the centralizer of a transitive group is semiregular.

Suppose that N_1 and N_2 are minimal normal subgroups of the primitive group G. Then each of N_1 and N_2 is transitive; but they centralize each other, and so each is semiregular, and so regular. Clearly it is not possible for there to be a third minimal normal subgroup.

The centralizer of a regular permutation group (in the symmetric group) is regular; indeed, the centralizer of the *right regular representation* of a group (acting on itself by right multiplication) is the *left regular representation*. These two regular groups coincide if and only if they are abelian. So, in our situation, N_1 and N_2 must be non-abelian. □

Example 2.6. Here is an example of a primitive group with two minimal normal subgroups. Let S be any finite group, and let $G = S \times S$. Then there is an action of G on S, where the first factor acts on the left and the second on the right, as follows:

$$(g, h) : x \mapsto g^{-1}xh.$$

It is easy to show that the action is faithful if and only if S has trivial centre. Moreover, the action is primitive if and only if S is a non-abelian simple group. For any congruence for the second factor is the relation "same coset of T" for some subgroup T of S; and this congruence is preserved by the first factor if and only if T is a normal subgroup. So, if S is simple, then G has only the trivial congruences; and conversely.

2.2.2. Other definitions. In this section, we give two further properties equivalent to primitivity, one due to Higman, the other to Rystsov.

Let G be a transitive permutation group on Ω. Then the set Ω^2 of ordered pairs of elements of Ω is partitioned into orbits under the componentwise action of G. These orbits are called *orbitals* of G. One orbital consists of all the pairs (a, a) for $a \in \Omega$ (by the assumption of transitivity); this is the *diagonal orbital*. Any non-diagonal orbital can be regarded as the set of edges of a digraph on the vertex set Ω, called an *orbital digraph* of G. If the orbital O is *symmetric* (that is, $(a, b) \in O$ implies $(b, a) \in O$), then we can regard the orbital digraph as an undirected graph.

Theorem 2.7. *The transitive permutation group G is primitive if and only if every non-diagonal orbital digraph is connected.*

Proof. If there is an orbital digraph which is not connected, then its connected components form blocks for G. Conversely, suppose that there is a non-trivial G-invariant partition P of Ω, and choose distinct points a, b in the same part of P. Then the orbital digraph corresponding to the orbital $(a, b)G$ has the property that all its edges are contained within parts of P, so it is not connected. □

The next theorem was essentially proved, but not stated, by Rystsov [91]; the statement in this form appears in [12].

Theorem 2.8. *Let G be a transitive permutation group on Ω, where $|\Omega| = n$. Then G is primitive if and only if, for any map $f : \Omega \to \Omega$ of rank $n - 1$, the monoid $\langle G, f\rangle$ contains an element of rank 1.*

This theorem is not difficult, but it will be much easier when we have developed a bit more technique.

2.3. Imprimitive groups and wreath products. A transitive but imprimitive permutation group G preserves a partition P of Ω, and so is contained in the group of all permutations fixing this partition. Since G is transitive, the partition is uniform, with (say) m parts each of size k. Then the group fixing the partition

is the *wreath product* S_k wr S_m of symmetric groups of degrees k and m. This means that it is the product of two subgroups:

(a) the *base group*, the direct product of m copies of S_k, where the ith copy acts on the ith part of the partition P;

(b) the *top group*, isomorphic to S_m, which permutes the parts of P.

The base group is a normal subgroup, and the top group acts on it by permuting the direct factors. Thus the group has order $(k!)^m\, m!$.

It is possible to define the wreath product of two arbitrary permutation groups similarly, and to show that if G is an imprimitive group, H is the group induced on a block by its setwise stabilizer, and L is the group of permutations of the blocks induced by G, then G is embedded into the wreath product H wr L.

The set on which the wreath product acts can be identified with the Cartesian product $K \times M$, where $K = \{1, \ldots, k\}$ and $M = \{1, \ldots, m\}$: we think of this as a fibre space over M, where each fibre is isomorphic to K (Figure 2):

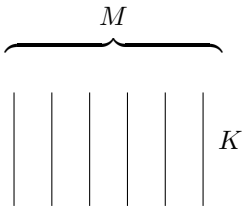

Figure 2. A fibration

This action of the wreath product is called the *imprimitive action*, as opposed to the *power action*, which we will meet shortly.

2.4. The O'Nan–Scott theorem. The group-theoretic structure of primitive groups was further elucidated independently by Michael O'Nan and Leonard Scott at a conference on finite groups in Santa Cruz in 1979. Both papers appeared in preliminary proceedings but only Scott's paper is in the final volume. We do not require the full detail of the theorem (for which see [78]), so we can make some simplifications.

The *socle* of a finite group is the product of its minimal normal subgroups. (Any two minimal normal subgroups commute, and each has trivial intersection with the product of the others, so we have a direct product.) As we saw above, a primitive group has at most two minimal normal subgroups, and if there are two then they are isomorphic; so the socle of a primitive group is a product of isomorphic finite simple groups. The O'Nan–Scott theorem allows us to apply the Classification of Finite Simple Groups to the study of primitive groups.

2.4.1. Non-basic groups. A *Cartesian structure* or *power structure* on Ω is a bijection between Ω and the set K^M of functions from M to K, where $|M|, |K| > 1$. This gives Ω the structure of an m-dimensional hypercube (where $m = |M|$) whose sides have size $|K|$. If $K = \{1, \ldots, k\}$ and $M = \{1, \ldots, m\}$, then Ω is identified with the set of m-tuples over the alphabet $K = \{1, \ldots, k\}$. The automorphism group of a power structure is the wreath product $S_k \operatorname{wr} S_m$, but in a different action from the imprimitive action we saw earlier: the *power action*, or *product action*, of the wreath product.

Let G act on Ω. We say that G is *non-basic* if it preserves a Cartesian structure on Ω, and *basic* otherwise.

A transitive non-basic group is embeddable in the wreath product of permutation groups on K and M in the power action. Elements of the base group of the wreath product permute the symbols in each coordinate independently, while elements of the top group permute the coordinates. If we take a set of size km partitioned into m sets of size k on which the wreath product has its imprimitive action, then we can identify the elements of the Cartesian structure with sets of points which are sections for the partition (that is, contain one element from each part).

In other language, in terms of the fibre space $K \times M$ on which the wreath product has its imprimitive action, the elements of the product action (which are the functions $\phi : M \to K$) are the *global sections* of the fibration (Figure 3):

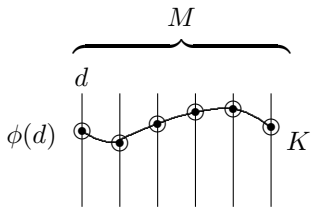

Figure 3. A global section

However, for primitive groups, we can make a stronger statement.

Theorem 2.9 (O'Nan–Scott for non-basic groups). *Let G be a primitive but non-basic permutation group with socle N. Then G is embeddable in the wreath product G_0 wr K, where G_0 is a basic primitive permutation group. Moreover, if K has degree n, then $N = N_0^n$, where N_0 is either the socle or a minimal normal subgroup of G_0.*

The case where G_0 has two minimal normal subgroups, of which N_0 is one (the so-called *twisted wreath product* case), was pointed out by Michael Aschbacher. The smallest twisted wreath product has degree $60^6 = 46656000000$. A discussion of these is given in [47]. It will turn out that non-basic groups are non-synchronizing, so we will not be concerned with twisted wreath products.

2.4.2. Basic groups. In order to describe basic groups, we need to look at several special classes of groups.

Affine groups Let V be a d-dimensional vector space over the field $\mathbb{F}_p$, where p is prime, and let H be a group of linear transformations of V. Then there is a corresponding *affine group*

$$G = \{x \mapsto xh + v : h \in H, v \in V\}$$

of permutations of V, generated by the translations (which form a normal subgroup) and elements of H.

Theorem 2.10. *With the above notation,*

(a) *G is always transitive;*

(b) *G is primitive if and only if H acts* irreducibly *on V (that is, leaves invariant no non-zero proper subspace of V);*

(c) *G is basic if and only if H acts* primitively *on V (that is, preserves no non-trivial direct sum decomposition of V).*

A primitive group is affine if and only if its socle (which is its unique minimal normal subgroup) is an elementary abelian p-group.

Diagonal groups Let S be a non-abelian finite simple group. A *diagonal group* is one whose socle is S^n, acting on the cosets of a diagonal subgroup

$$\{(s, s, \ldots, s) : s \in S\}$$

of S^n.

For $n = 2$ we have the example of $S \times S$ acting by left and right multiplication we saw earlier.

A diagonal group may also contain

(a) automorphisms of S, acting in the same way on all factors;

(b) permutations of the factors.

If $n > 2$, we must have at least a transitive group of permutations of the factors in order for the diagonal group to be primitive.

Almost simple groups A group G is *almost simple* if its socle is simple. Such a group is an extension of a simple group by a subgroup of its automorphism group; in other words, there is a simple group S such that $S \leq G \leq \mathrm{Aut}(S)$.

For example, the symmetric group S_n is almost simple for $n \geq 5$. (It is affine for $n \leq 4$.)

The almost simple primitive groups are the largest and least understood class. Note that, unlike the other two cases, the action of the group is not specified.

Theorem 2.11 (O'Nan–Scott for basic groups). *Let G be a basic primitive permutation group. Then G is affine, or diagonal, or almost simple.*

The O'Nan–Scott Theorem opened the way to the application of the Classification of Finite Simple Groups to permutation group theory, which has been done very successfully since the Classification was first announced in 1980. (These results were conditional on the Classification until its proof was completed in 2005.)

2.5. The Classification of Finite Simple Groups. This major theorem has a proof which is currently over 10000 pages long. We will not specify the groups too precisely, since there are good sources for this: we recommend [115].

Theorem 2.12. *Any finite simple group is one of the following:*

(a) *a cyclic group of prime order;*

(b) *an alternating group A_n, $n \geq 5$;*

(c) *a group of Lie type;*

(d) *one of the* 26 *sporadic finite simple groups.*

We refer to this theorem as CFSG.

The groups of Lie type are quotients of matrix groups over finite fields. There are finitely many families; some (the *classical groups*, which we will discuss in more detail later) are parametrized by a dimension and a field order; the rest (the *exceptional groups*) just by a field order.

2.6. 2-transitive and 2-homogeneous groups. A permutation group G on Ω is said to be 2-*transitive* if it acts transitively on the set of ordered pairs of distinct elements of Ω: in other words, given two ordered pairs (a_1, a_2) and (b_1, b_2), with $a_1 \neq a_2$ and $b_1 \neq b_2$, there exists $g \in G$ with $a_i g = b_i$ for $i = 1, 2$.

A permutation group G on Ω is said to be 2-*homogeneous* if it acts transitively on the set of 2-element subsets of Ω. This is weaker than 2-transitivity, since we do not require that we can interchange two points. Indeed, a 2-homogeneous group is 2-transitive if and only if its order is even. For, if G is 2-transitive, then an element which interchanges two points has even order. Conversely, a group of even order contains an involution, so some pair of points can be interchanged; if the group is 2-homogeneous, then any pair can be interchanged.

Using this, the classification of 2-homogeneous but not 2-transitive groups was achieved by Kantor and Berggren independently in the late 1960s, using the Feit–Thompson theorem on solvability of groups of odd order [69, 22].

Theorem 2.13. *Let G be a permutation group on Ω which is 2-homogeneous but not 2-transitive. Then we can identify Ω with the finite field $\mathbb{F}_q$ where $q \equiv 3 \pmod 4$, so that G is a subgroup of the semi-affine group*

$$\{x \mapsto a^2 x^\sigma + b : a, b \in \mathbb{F}_q, a \neq 0, \sigma \in \mathrm{Aut}(\mathbb{F}_q\}.$$

Proof. The group G has odd order, so by the Feit–Thompson theorem it is solvable. Hence it has an elementary abelian regular normal subgroup N which is the additive group of a vector space. Now consider the group $G^+ = \langle G, -1\rangle$. This group is 2-transitive, and also solvable. The 2-transitive solvable groups were determined by Huppert [65]; by examining the list we can complete the proof. □

The classification of 2-transitive groups is a consequence of CFSG. We do not discuss the result here; the list of 2-transitive groups can be found in the books by Cameron [35] and by Dixon and Mortimer [47].

We conclude this section with a simple observation:

Proposition 2.14. *A 2-homogeneous group is primitive.*

Proof. This follows easily from Higman's Theorem, since any two points are adjacent in a non-trivial orbital graph for G, which is thus connected. □

So we have the following properties of permutation groups:

$$\begin{aligned}\text{transitive} &\Leftarrow \text{primitive} \Leftarrow \text{basic}\\ &\Leftarrow \text{2-homogeneous} \Leftarrow \text{2-transitive.}\end{aligned}$$

Note that each of these properties is closed upwards: an overgroup of a permutation group with the property also has the property.

We will extend this hierarchy by inserting some new classes of permutation groups between 2-homogeneous and basic.

3. Synchronization

Our automata (illustrated by examples below) are very simple gadgets. An *automaton* is a triple (Ω, A, F), where Ω is a finite set of states, A a finite input alphabet, and F a map from A to the set of all functions from Ω to itself. The machine is in one of the states $s \in \Omega$; on reading an *input*, a symbol from the alphabet A, it undergoes a change of state to the image of s under $F(a)$. Unlike the automata used in other areas, there is no prescribed initial or terminal state; these automata do not accept languages. Moreover, they are deterministic: there is a unique function $F(a)$ for each $a \in A$.

An automaton can read a *word* in A, a finite sequence of letters; each letter causes a change of state, and the overall effect is then a function on Ω which is the composition of the functions corresponding to the letters. Our interest is whether an automaton has a *reset word* with the property that, if it is presented in an unknown state, then reading the reset word brings it to a known state.

3.1. Two examples. The first example was suggested by Olof Sisask.

Example 3.1. A certain calculator has an 'On' button but no 'Off' button. To switch it off, you hold down the 'Shift' key and press the 'On' button. The 'Shift' key has no effect if the calculator is switched off. Assuming that you can't see the screen, how can you ensure that the calculator is switched off?

Obviously, pressing the 'On' button leaves the calculator switched on, no matter what its former state; and then 'Shift-On' will switch it off.

Note that if, instead, there is a single 'On-Off' button which toggles the states, then the problem would have no answer.

Example 3.2. You are in a dungeon consisting of a number of rooms. Passages are marked with coloured arrows, where the colours are *dotted* and *solid*. Each room contains a special door; in one room (say 4), the door leads to freedom, but in all the others, to instant death. You have a schematic map of the dungeon (Figure 4), but you do not know where you are.

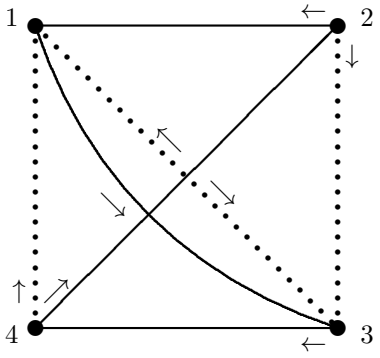

Figure 4. An automaton

You can check that the sequence (dotted, solid, dotted) takes you to room 1 from any starting point. Then you can use the map to navigate to the exit door in room 4.

3.2. Automata and synchronization. Combinatorially, an automaton can be regarded as an edge-coloured digraph, where the vertices are the states, and edge colours correspond to letters of the alphabet; there should be one edge with each colour out of each vertex. On reading a letter in a given state, the automaton follows the edge whose colour corresponds to that letter, and moves to the state at the terminus of this edge. This is the representation used in Figure 4. A n-coloured automaton is said to be *complete and deterministic* if for each colour c and from each vertex (or state) there is one and only one out-arrow coloured c. In this paper we only consider deterministic and complete automata. Figure 4 provides the example of a 4-state 2-coloured automaton which is complete (in every state there are out-arrows of both colours) and deterministic (no state has two different out-arrows of the same colour). In a finite deterministic automaton, each colour induces a transformation of the states. For example, in Figure 4 the two colours, that is, the two transitions, induce the following transformations

$$\text{dotted} = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 3 & 3 & 1 & 1 \end{pmatrix} \quad \text{solid} = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 3 & 1 & 4 & 2 \end{pmatrix}.$$

Algebraically, we will be interested in the composition of transitions of the automaton. The set of all transformations of the states which are produced by applying a (possibly empty) sequence of transitions is a *transformation semigroup* on Ω; that is, a set of transformations closed under composition. So we can regard an automaton as a transformation semigroup (acting on the set of states) with a prescribed set of generators (the transitions of the automaton, corresponding to the colours or the letters labelling the arrows).

The *rank* of a transformation of Ω is the cardinality of its image.

A *reset word* is a sequence of transitions such that the composition of the transitions in the sequence, applied to any starting vertex, brings you to the same state. An automaton which possesses a reset word is called *synchronizing*. Thus, from the algebraic point of view, an automaton is synchronizing if the corresponding transformation monoid contains a map of rank 1, that is, a constant map.

Not every finite automaton has a reset word. For example, if every transition is a permutation, then every word in the transitions is a permutation, and has rank equal to $|\Omega|$.

3.3. The Černý conjecture. Here is a simplified example of the application of synchronization in industrial robotics (cf. [50]). The general situation is as follows.

Pieces arrive to be assembled by a robot. The orientation is critical. You could equip the robot with vision sensors and manipulators so that it can rotate the pieces into the correct orientation. But it is much cheaper and less error-prone to regard the possible orientations of the pieces as states of an automaton on which

transitions can be performed by simple machinery, and apply a reset word before the pieces arrive at the robot.

Example 3.3. Suppose that the component is square with a projecting tab on one side.

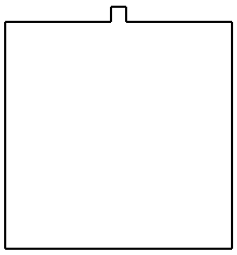

It can sit in a tray on the conveyor belt in any one of four orientations.

The following transitions are easy to implement:

A: rotate through 90° in the positive direction;

B: rotate through 90° if the projection points up, otherwise do nothing.

Figure 5 is a diagram of the automaton. Transition A is represented by a dotted line; transition B is represented by a solid line (fixing the states 2, 3 and 4). Each state represents the position of the component with the projection on that side.

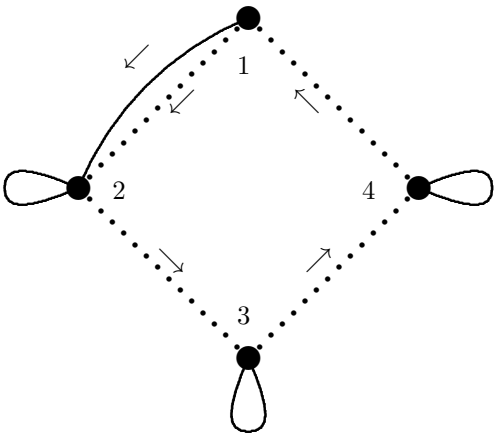

Figure 5. Another automaton

Now the following table is easily checked.

	B	A	A	A	B	A	A	A	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 BAAABAAAB is a reset word.

It can be shown that there is no shorter reset word for this automaton.

Moreover, the example extends to any number n of states, replacing the square by a regular n-gon. The corresponding shortest reset word has length $(n-1)^2$.

In 1969, Černý made the following conjecture (see [108]):

Conjecture 3.4. Suppose that an automaton with n states is synchronizing. Then it has a reset word of length at most $(n-1)^2$.

This conjecture is still open after close to fifty years! The example above shows that, if true, it would be best possible. The best current upper bound, despite years of intensive effort, is $\frac{n^3-n}{6}$, due to the combined work of Frankl [54] and Pin [87] from 1983. The remainder of the literature consists of special cases (e.g. [4, 5, 6, 7, 9, 14, 21, 40, 41, 42, 50, 70, 71, 86, 88, 89, 90, 91, 92, 97, 99, 100, 104, 105]). The strongest result is Dubuc's theorem [49], which proves the Černý conjecture under the assumption that some transition cyclically permutes the states, as is the case for the Černý examples [42]. See [68] for a recent approach involving linear programming.

One of the difficulties of the Černý problem is that there are few known families of slowly synchronizing automata (cf. [8] for the connection with exponents of primitive digraphs) and so we don't have a very good understanding of what makes an automaton slow to synchronize. In fact, the Černý sequence is still the only infinite sequence of examples of n-state automata that have minimal length reset word of length $(n-1)^2$.

The other issue is that random automata are synchronizing and synchronize quickly. More precisely, Berlinkov [23] showed that a random n-state automaton is synchronizing with high probability as n approaches infinity. Nicaud [85] proved that if $\varepsilon > 0$, then with high probability an n-state automaton has a reset word of length at most $n^{1+\varepsilon}$. Thus one is unlikely to find a counterexample by searching at random; also the search space is too large for a brute force attempt to find a counterexample.

We will not prove the Černý conjecture in this paper, but it provided motivation for our approach, and we will return to it later.

For an accessible discussion of the Černý conjecture, we recommend the survey by Volkov [108].

3.4. The Road-Colouring Conjecture. The underlying digraph of an automaton with n transitions is a digraph with the property that every vertex has exactly n arcs leaving it.

Conversely, and trivially, given any digraph with this property, it is clear that it can be edge-coloured so as to represent an automaton.

The resulting automaton may or may not be synchronizing. What are necessary and sufficient conditions for there to be an edge-colouring representing a synchronizing automaton?

We will assume that the automaton can be synchronized in any given state by a suitable reset word. A necessary and sufficient condition for this is that it is strongly connected. (If so then, as in our dungeon, if we can synchronize at some state, we can move from there to any other state.)

It is also necessary that the greatest common divisor of the lengths of cycles in the digraph is 1. For suppose the g.c.d. of cycle lengths is d. Choose any vertex v, and let Ω_i be the set of vertices reachable from v in a number of steps congruent to i mod d, for $i = 0, 1, \ldots, d-1$. The sets Ω_i are pairwise disjoint, and so no automaton based on the digraph can be synchronizing.

The conjecture that these two necessary conditions are also sufficient was made in 1970 by Weiss and Adler [1, 2] in connection with symbolic dynamics, and became known as the *Road-Colouring Conjecture*. It was proved by Avraham Trahtman in 2007 [106]:

Theorem 3.5. *Let D be a digraph which is strongly connected and has constant out-degree, and suppose that the greatest common divisor of the cycle lengths in D is 1. Then D can be edge-coloured so as to produce a synchronizing automaton.*

3.5. Synchronizing groups. Looking at the extreme examples above for the Černý conjecture, we see that, of the two transitions, the first is a cyclic permutation, which generates a transitive group on the set of states; the second is a non-permutation.

This observation is the basis of the next definition [10, 14]. A permutation group G on Ω is said to be *synchronizing* if, whenever f is a map on Ω which is not a permutation, the monoid $\langle G, f\rangle$ is synchronizing (that is, there is a word in f and the elements of G which has rank 1).

We have abused language here since G itself (regarded as a monoid) is not a synchronizing monoid; but a permutation group cannot be a synchronizing monoid, so hopefully the confusion will not be too great.

For example, the automorphism group of the Petersen graph is synchronizing. This fact can be proved by considering all possible non-permutations on the vertex set; but in the next section we will develop a technique to make it much easier to check assertions like this.

More generally, we say that a permutation group G *synchronizes* a non-permutation f if $\langle G, f\rangle$ contains a map of rank 1. Thus, G is synchronizing if it synchronizes every non-permutation.

The next theorem shows how synchronizing groups relate to the classical notions of primitive and 2-homogeneous groups.

Theorem 3.6. (a) *A synchronizing group is primitive.*

(b) *A 2-homogeneous group is synchronizing.*

Proof. (a) The simplest argument here is to recall the characterization of primitivity based on Theorem 2.8: a permutation group G of degree n is primitive if and only if the monoid $\langle G, f\rangle$ is synchronizing for any map f of rank $n-1$.

Since we haven't proved this yet, we give a different proof. Suppose that G is imprimitive. Let P be a non-trivial G-invariant partition of Ω, and let A be a subset of Ω containing one point from each part of P (so A is a *section*, or *transversal*, of P.) Now the map f that takes each point of Ω to the unique point of A in the same part of P is not synchronized by G. For any word in f and the elements of G which contains f at least once has the property that its image is a section for P, so no such word can have rank 1.

(b) Suppose that G is 2-homogeneous, and let f be any non-permutation. Let r be the minimal rank of an element $h \in \langle G, f\rangle$, and suppose for a contradiction that $r > 1$. Choose two distinct points x, y in the image of h, and two points u, v which are mapped to the same place by f. Then choose $g \in G$ mapping $\{x, y\}$ to $\{u, v\}$. Then hgf has smaller rank than h, a contradiction. So $r = 1$, as required. □

The first part of this theorem can be improved:

Proposition 3.7. *A synchronizing group is basic.*

Proof. Let G be non-basic, and suppose that Ω has been identified with the set of m-tuples over a set A of size k, in such a way that G preserves the identification (and so is embedded in S_k wr S_m).

Let f be the map which takes the m-tuple $(a_1, a_2, \ldots, a_m)$ to the m-tuple $(a_1, a_1, \ldots, a_1)$ with all entries equal. Let B be the image of f. Then applying any element of G to B gives a set of k elements whose projections onto any coordinate form the whole of A; so following this by f gives us the set B again. So no word in f and G can have rank smaller than k, and G fails to synchronize f. □

Figure 6 shows how the map works for S_3 wr S_2, the automorphism group of the 3×3 grid.

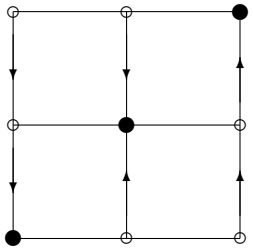

Figure 6. Failure to synchronize a square grid

We conclude this section with examples to show that the inclusions just proved do not reverse: we give examples of basic primitive groups which are not synchronizing, and synchronizing groups which are not 2-homogeneous.

Our examples are given by the symmetric groups S_m for $m \geq 5$, in their action on the set of 2-element subsets of $\{1, \ldots, m\}$.

(a) The group is primitive and basic for all $m \geq 5$. For it is easy to see that there are just two complementary orbital graphs: the vertex set is the set of 2-subsets of $\{1, \ldots, m\}$: two vertices are joined in the first graph if they have non-empty intersection (so this is the *line graph* of the complete graph K_m), and in the second if they have empty intersection. Now both of these graphs are connected.

(b) The group is not 2-homogeneous. For the edges in the two orbital graphs are not equivalent.

(c) The group is synchronizing if and only if m is odd. We will defer the proof of this assertion to the next section, when we will have another technique available.

3.6. Section-regular partitions. In the next section we will develop a very convenient combinatorial characterization of synchronization. In the meantime we give another characterization which was introduced in [10] and developed by Peter Neumann [84].

Let P be a partition of Ω. A *section*, or *transversal*, for P is a subset A of Ω which meets every part of P in a single point. Recall that a partition P is *non-trivial* if it is not the partition into singletons and not the partition with a single part.

Now let G be a permutation group on Ω. We say that the partition P is *section-regular* for G, with section A, if Ag is a section for P for every $g \in G$.

In Figure 6, the partition into vertical lines is section-regular for the group S_3 wr S_2 of automorphisms, with the diagonal as a section.

Theorem 3.8. *A permutation group G on Ω is synchronizing if and only if it has no non-trivial section-regular partition.*

Proof. Suppose first that G is non-synchronizing. Let f be a map such that $\langle G, f\rangle$ contains no map of rank 1. We may suppose, without loss of generality, that f is an element of minimal rank (say r) in $\langle G, f\rangle$. Let P be the kernel of f, and A the image of f. If Ag is not a section for P, then Ag meets fewer than r kernel classes of f, and so fgf has rank smaller than r, a contradiction. So P is section-regular with section A.

Conversely, suppose that P is a section-regular partition for G, with section A. Then for any $x \in \Omega$, there is a unique $y \in A$ which lies in the same class of P as does x. Define a map f by the rule that $xf = y$ when the above holds. Now Ag is a section for P, so $Agf = A$, for any $g \in G$. An easy induction shows that no element of $\langle G, f\rangle$ has rank smaller than $|A|$. □

The next two results are due to Peter Neumann [84]. A partition is said to be *uniform* if all its parts have the same size.

Theorem 3.9. *A section-regular partition for a transitive permutation group G on Ω is uniform.*

Proof. Let $n = |\Omega|$ be the degree of G. Suppose that P is section-regular for G, with section A. Suppose that B is a part of P. Count triples (a, b, g), where $a \in A$, $b \in B$ and $g \in G$ satisfies $ag = b$; there are $|A| \cdot |B|$ choices of a and b and then $|G|/n$ choices of g (since the set of elements of G mapping a to b is a coset of the stabilizer of a). On the other hand, for every $g \in G$, $|Ag \cap B| = 1$, so there is a unique pair (a, b) satisfying the condition. Thus $|A| \cdot |B| = n$; in particular, $|B|$ is independent of the part B of P chosen. □

Corollary 3.10. *A map f of minimal rank subject to being not synchronized by the transitive group G has uniform kernel.*

Proof. This follows from the two preceding theorems and the argument in the proof of Theoroem 3.8. □

It follows that any transitive group of prime degree is synchronizing, a result due originally to Pin [86] that can be considered the first result in the theory of synchronizing groups.

Theorem 3.9 and the proof of Theorem 3.8 in fact yield the following corollary.

Corollary 3.11. *Let M be a transformation monoid on a finite set Ω with a transitive group of units. Then each element of M of minimal rank has a uniform kernel.*

3.7. The Černý conjecture revisited. Can we prove at least some instances of the Černý bound for transformation monoids of the form $\langle G, f\rangle$, where G is a synchronizing permutation group?

Since the permutations in this monoid are precisely the elements in G, and these by themselves will not synchronize, it seems reasonable to build a word of the form $fg_1fg_2 \cdots fg_rf$, where we use f to reduce the rank of the partial product and g_i to ensure that the next application of f does so. Note that the rank of hf is strictly less than the rank of f if and only if two points of the image of h lie in the same kernel class of f. So, if the rank of a product of the above form is at most k, choose g_{r+1} to map two points in the image of the product into a kernel class of f, and then the rank of the product $fg_1f \cdots g_{r+1}f$ will be at most $k - 1$.

If this strategy succeeds, we will have a reset word with at most $n - 1$ occurrences of f. The task now is to bound the lengths of the expressions for $g_1, \ldots$ in terms of the given generators of G, for which hopefully group theory will help.

We want to avoid the case where there is a set A (the image of a subword of the product), all of whose G-images are partial sections for the kernel of f. This is where conditions on G like "synchronizing" are relevant.

4. Graph endomorphisms

In this section, we define endomorphisms of graphs, and use them to give characterizations of synchronizing monoids and groups. This result gives us the simplest available test for the synchronizing property of permutation groups. We will illustrate by returning to the example of the symmetric group S_m acting on 2-sets, and showing that it is synchronizing if and only if m is odd.

4.1. Cliques, colourings and endomorphisms. A *relational structure* consists of a set carrying a number of relations of specified arities. The most important example for us is a graph, a set with a single binary relation.

A *homomorphism* $f : A \to B$ between relational structures A and B is a map between the underlying sets which preserves all instances of the relation. Thus, a graph homomorphism maps edges to edges, but its action on non-edges is not specified; it could map a non-edge to an edge, or to a non-edge, or to a single vertex. (If the graph has no loops, then edges cannot be collapsed to single vertices.)

Some important graph parameters can be expressed in terms of homomorphisms. The *clique number* $\omega(\Gamma)$ of a graph Γ is the number of vertices in the largest complete subgraph of Γ, that is, the largest number of vertices such that any two are adjacent. A *(proper) colouring* of Γ is an assignment of colours from a set C to the vertices in such a way that the ends of any edge have different colours. The *chromatic number* $\chi(\Gamma)$ of Γ is the smallest number of colours required for a proper colouring of Γ.

It is clear that $\chi(\Gamma) \geq \omega(\Gamma)$ for any graph Γ, since the vertices of a clique must all have different colours.

Recall that K_r is the complete graph with r vertices.

Theorem 4.1. (a) *There is a homomorphism from K_r to Γ if and only if $\omega(\Gamma) \geq r$.*

(b) *There is a homomorphism from Γ to K_r if and only if $\chi(\Gamma) \leq r$.*

Proof. (a) The images of the vertices of K_r under a homomorphism must all be distinct.

(b) Let C be the vertex set of K_r. We think of C as a set of colours, and the homomorphism f assigns to v the colour $f(v)$. Now the definition of a homomorphism shows that this is a proper colouring. □

Corollary 4.2. *For a graph Γ, the following are equivalent:*

(a) $\omega(\Gamma) = \chi(\Gamma)$*;*

(b) *there are homomorphisms in both directions between Γ and a complete graph.*

For a detailed study of graph homomorphisms, we recommend [59].

An *endomorphism* of Γ is a homomorphism from Γ to itself. The composition of endomorphisms is an endomorphism, and the identity map is an endomorphism; so the set of endomorphisms of Γ is a transformation monoid on the vertex set of Γ, denoted by $\mathrm{End}(\Gamma)$.

In line with our previous practice, we write endomorphisms on the right.

4.2. Graphs and endomorphism monoids. The map $\Gamma \to \mathrm{End}(\Gamma)$ is a mapping from graphs to transformation monoids. Unfortunately, it is not a functor in any reasonable sense; this is also the case for the next map we define, which goes in the other direction.

Let M be a transformation monoid on a set Ω. We define a graph $\Gamma = \mathrm{Gr}(M)$ on the vertex set Ω by the following rule for adjacencies:

> v and w are adjacent in $\mathrm{Gr}(M)$ if and only if there does not exist $f \in M$ with $vf = wf$.

This correspondence has various nice properties:

Theorem 4.3. (a) *For any transformation monoid M on Ω, the graph $\mathrm{Gr}(M)$ has the properties*

(i) $M \leq \mathrm{End}(\mathrm{Gr}(M))$;

(ii) $\omega(\mathrm{Gr}(M)) = \chi(\mathrm{Gr}(M))$.

(b) *If $M_1 \leq M_2$, then $\mathrm{Gr}(M_2)$ is a spanning subgraph of $\mathrm{Gr}(M_1)$.*

Proof. (a) (i) Let $f \in M$; we have to show that f is an endomorphism of $\mathrm{Gr}(M)$, so suppose not. Then there exists an edge $\{v, w\}$ of $\mathrm{Gr}(M)$ which is not preserved by M. By definition, $vf \neq wf$; so this can only happen if $\{vf, wf\}$ is a non-edge of $\mathrm{Gr}(M)$. But then, by definition, there exists $h \in M$ such that $(vf)h = (wf)h$. Then $fh \in M$ and $v(fh) = w(fh)$, contradicting the fact that $\{v, w\}$ is an edge of $\mathrm{Gr}(M)$.

(a) (ii) Now let f be an element of M of smallest possible rank. Let $A = \mathrm{Im}(f)$. No element of M can map two points of A to the same place, since if h did so then fh would have smaller rank than f. So by definition, A is a clique in $\mathrm{Gr}(M)$. Since $f \in \mathrm{End}(\mathrm{Gr}(M))$, we see that f induces a proper colouring of $\mathrm{Gr}(M)$ with $|A|$ colours.

(b) Clearly adding extra endomorphisms cannot produce new edges which were not there before. □

4.3. Characterization of synchronizing monoids. Now we can give our characterization of synchronizing monoids.

Theorem 4.4. *Let M be a transformation monoid on Ω. Then M is non-synchronizing if and only if there exists a non-null graph Γ on the vertex set Ω with $M \leq \mathrm{End}(\Gamma)$. If such a graph Γ exists, then we may choose it so that $\omega(\Gamma) = \chi(\Gamma)$, and this number is equal to the minimum rank of an element of M.*

Proof. If $M \leq \mathrm{End}(\Gamma)$ for any non-null graph Γ, then M is not synchronizing, since edges of Γ cannot be collapsed by elements of M.

Conversely, if M is non-synchronizing, let $\Gamma = \mathrm{Gr}(M)$. Suppose that Γ is the null graph. Then any pair of points of Ω are mapped to the same place by some element of M. Let f be an element of least possible rank in M. If the rank of f is greater than 1, choose $x, y \in \mathrm{Im}(f)$, and $h \in M$ with $xh = yh$; then fh has

smaller rank than f. So f has rank 1, and M is synchronizing, a contradiction. So Γ is non-null.

Now the remaining assertions of the theorem come from the properties of $\mathrm{Gr}(M)$ from the preceding subsection. □

Corollary 4.5. *Let G be a transitive permutation group on Ω. Then G is synchronizing if and only if every non-trivial G-invariant graph Γ has $\omega(\Gamma) \neq \chi(\Gamma)$.*

Proof. If Γ is a non-trivial G-invariant graph with $\omega(\Gamma) = \chi(\Gamma) = r$, then there are graph homomorphisms $f\colon \Gamma \to K_r$ and $g\colon K_r \to \Gamma$. Composition of these homomorphisms provides a singular endomorphism $h\colon \Gamma \to \Gamma$ and so $\langle G \cup h\rangle$ is not synchronizing by the previous theorem.

The converse follows immediately from Theorem 4.4. □

For example, the automorphism group of the Petersen graph is edge-transitive and nonedge-transitive; so we only have to check the Petersen graph and its complement. It is not hard to show that the Petersen graph has clique number 2 and chromatic number 3, while its complement has clique number 4 and chromatic number 5. So the automorphism group is synchronizing.

This corollary is the basis for the best computational test for synchronization. The test runs as follows. Given a transitive permutation group G on Ω, run the following algorithm:

Algorithm 4.6. (a) Find all the non-trivial G-invariant graphs. There are $2^r - 2$ such graphs, where r is the number of orbits of G on 2-element subsets of Ω, since the edge set of a G-invariant graph is a union of orbits of G.

(b) Test each graph Γ to see whether $\omega(\Gamma) = \chi(\Gamma)$. If one does, then G is not synchronizing; otherwise it is synchronizing.

This algorithm looks extremely inefficient. The first stage generates exponentially many graphs to be checked; and computing the clique number and chromatic number of a graph are both NP-complete problems.

However, in practice, "interesting" permutation groups often have comparatively few orbits on 2-sets, so r is small; and the graphs which have to be tested have large automorphism groups, which can be exploited to reduce the computational burden in the second step.

In the next section, we will see how the algorithm can be slightly improved.

Example 4.7. Here is an example promised earlier. Let G be the symmetric group of degree $m \geq 5$, in its action on 2-element subsets of $\{1, \ldots, m\}$. There are only two orbits on pairs of 2-element subsets: the subsets may intersect in a point, or they may be disjoint. So we have two G-invariant graphs to consider: the line graph of K_m and its complement.

Let Γ be the line graph of K_m. The clique number of Γ is $m-1$; a typical maximal clique is $\{\{1, i\} : i = 2, \ldots, m\}$. When can the chromatic number be $m-1$? Pairs with the same colour must be disjoint, so there are at most $\lfloor m/2 \rfloor$ pairs in a colour class. If m is odd, this number is $(m-1)/2$, so at least $\binom{m}{2}/((m-1)/2) = m$

colours are required. If m is even, we can have $m/2$ edges in a colour class, and $m-1$ colours. This can be realized as follows. Take a regular $(m-1)$-gon in the plane. The edges and diagonals fall into $m-1$ parallel classes, with $(m/2)-1$ pairs in each class, and one point omitted from each class. Assign one colour to all the edges in each class. Now add an extra point ∞, and give the colour of a class C to the pair consisting of ∞ and the point omitted by C. For example, if $m = 6$ and we label the vertices of the regular pentagon by $1, 2, 3, 4, 5$ in counterclockwise ordering, then the five colour classes are $\{\{1,2\},\{3,5\},\{4,\infty\}\}$, $\{\{2,3\},\{1,4\},\{5,\infty\}\}$, $\{\{3,4\},\{5,2\},\{1,\infty\}\}$, $\{\{4,5\},\{1,3\},\{2,\infty\}\}$, and $\{\{1,5\},\{2,4\},\{3,\infty\}$. So this graph Γ has $\omega(\Gamma) = \chi(\Gamma)$ if and only if m is even.

Now let Γ be the complement of the line graph of K_m. Now a clique consists of disjoint pairs, so the largest clique has size $\lfloor m/2 \rfloor$. But Γ cannot be coloured with this many colours. For the colour classes must be cliques in $L(K_m)$; we saw that such cliques have size at most $m-1$ and so we would need at least $m/2$ in a partition. So we could only achieve the bound if m were even and the cliques were pairwise disjoint. But the cliques of size $m-1$ consist of all pairs containing a given point; and the cliques defined by points a and b have the pair $\{a,b\}$ in common. So this graph never has clique number and chromatic number equal.

We conclude that, for $m \geq 5$, S_m acting on 2-sets is synchronizing if and only if m is odd.

We remark that, in fact, the chromatic number of the complement of $L(K_m)$ is known to be $m-2$; this is a special case of a theorem of Lovász [80].

Theorem 4.4 and Corollary 4.5 have been used in a number of places, for example, [11, 12, 37, 93], for investigating synchronizing groups.

4.4. Rystsov's Theorem. We illustrate these concepts by proving Theorem 2.8. With the terminology we have introduced, the theorem states:

Theorem 4.8. *A transitive permutation group of degree n is primitive if and only if it synchronizes every map of rank $n-1$.*

Proof. Suppose first that G fails to synchronize the map f of rank $n-1$. Then there exist a and b such that $af = bf$, but f is injective on any subset not containing both a and b. Suppose that Γ is a non-trivial graph with $\langle G, f\rangle \leq \mathrm{End}(\Gamma)$. Since G is transitive, Γ is regular; suppose that every vertex has degree d. Since $af = bf$, we see that $\{a,b\}$ is a non-edge of Γ; so f maps the neighbours of a bijectively to the neighbours of af. Similarly, f maps the neighbours of b bijectively to the neighbours of $bf = af$. Hence a and b have the same neighbours. Now the relation $\equiv$, defined by $x \equiv y$ if and only if x and y have the same neighbours, is a G-invariant equivalence relation; so G is imprimitive.

Conversely, suppose that G is imprimitive; let P be a non-trivial G-invariant partition. Let a and b be two points in the same part of P. Define a map f by

$$xf = \begin{cases} x & \text{if } x \neq b; \\ a & \text{if } x = b. \end{cases}$$

It is easy to see that f has rank $n-1$ and is not synchronized by G (it is an endomorphism of the complete multipartite graph with multipartition P). □

In the paper [11], the authors extend this result to show that a primitive group of degree n synchronizes any map of rank $n-4$ or greater. The non-basic group S_3 wr S_2, the automorphism group of the 3×3 square grid, has degree 9 and fails to synchronize a map of rank 3 (the grid graph has clique number and chromatic number 3); so this result is within one of best possible.

4.5. Cores and hulls. The *core* of a graph Γ is the smallest graph Δ with the property that there are homomorphisms from Γ to Δ and from Δ to Γ. It is known that every graph has a core, which is unique up to isomorphism; moreover, the core is an induced subgraph of Γ, and there is a *retraction* from Γ to its core (an endomorphism which acts as the identity on its image).

Cores play an important role in the theory of graph homomorphisms, see [59]. We remark that the graphs which we used in our characterization of synchronizing monoids can be defined as the graphs whose cores are complete. The following well-known result in graph theory [58, Theorem 3.9], can be viewed as part of the theory of synchronizing groups.

Theorem 4.9. *Let Γ be a vertex-transitive graph. Then the retraction from Γ to its core is uniform; in particular, the number of vertices in* $\mathrm{Core}(\Gamma)$ *divides the number of vertices of Γ.*

Proof. The retraction from Γ to its core is a minimal rank element of $\mathrm{End}(\Gamma)$ and so the theorem follows from Corollary 3.11. □

A "dual" concept is that of the *hull* of a graph, introduced in [38]. It is defined by

$$\mathrm{Hull}(\Gamma) = \mathrm{Gr}(\mathrm{End}(\Gamma));$$

in other words, two vertices are adjacent in the hull of Γ if and only if no endomorphism of Γ collapses them to the same point.

Some of its properties are given by the following result:

Theorem 4.10. (a) *Γ is a spanning subgraph of* $\mathrm{Hull}(\Gamma)$.

(b) *The core of* $\mathrm{Hull}(\Gamma)$ *is a complete graph on the vertices of the core of Γ.*

(c) $\mathrm{End}(\Gamma) \leq \mathrm{End}(\mathrm{Hull}(\Gamma))$ *and* $\mathrm{Aut}(\Gamma) \leq \mathrm{Aut}(\mathrm{Hull}(\Gamma))$.

Proof. (a) This just says that every edge of Γ is an edge of its hull, which is clear since endomorphisms do not collapse edges.

(b) The core of Γ is the image of an endomorphism of Γ of minimal rank. Thus endomorphisms of Γ cannot identify two vertices of the core, so it induces a clique in $\mathrm{Hull}(\Gamma)$. This clique is the image of an endomorphism of $\mathrm{Hull}(\Gamma)$, and it is clear that no endomorphism can have smaller image; so it is the core of $\mathrm{Hull}(\Gamma)$.

(c) Putting $M = \mathrm{End}(\Gamma)$, we know that $M \leq \mathrm{End}(\mathrm{Gr}(M))$, which gives the first inequality; the second follows. □

Thus, passing from a graph to its hull cannot decrease the symmetry, but might increase it in some cases.

Example 4.11. Let Γ be the path of length 3, shown in Figure 7.

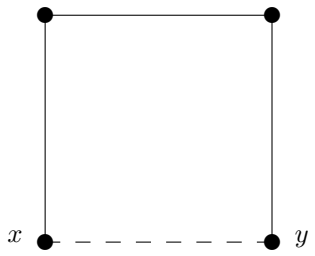

Figure 7. A graph and its hull

No homomorphism can identify x and y, so they are joined in the hull. Note the increase in symmetry: $|\operatorname{Aut}(\Gamma)| = 2$ but $|\operatorname{Aut}(\operatorname{Hull}(\Gamma))| = 8$.

5. Related concepts

The definition of synchronization can be varied in several ways, giving rise to several closely-related concepts. We consider some of these in this section of the paper. Perhaps the most interesting is the property of being almost synchronizing. The first examples showing that this is not equivalent to primitivity were found very recently.

We also define some measures of how far a given group is from being synchronizing.

5.1. Almost synchronizing groups. We saw in the preceding section that the symmetric group S_m acting on 2-sets is primitive but not synchronizing if m is even and $m \geq 6$.

However, the only maps that it fails to synchronize are the proper endomorphisms of $L(K_m)$, which all have rank $m-1$.

To take another example, consider the (non-basic primitive) group S_k wr S_m for $k \geq 3$, $m \geq 2$. This group fails to be synchronizing, but the situation is more complicated than the preceding one. Consider the *Hamming graph* $H(m,k)$, whose vertices are all m-tuples over an alphabet of size k, and two m-tuples are adjacent if they agree in all but one position. The automorphism group of this graph is the wreath product S_k wr S_m.

Now, for any d with $1 \leq d \leq m$, we can construct an endomorphism f_d with rank k^d, as follows. Choose the alphabet to be a group of order k, for example, the additive group of the integers mod k. Now set

$$(a_1, \ldots, a_m)f_d = (a_1, \ldots, a_{d-1}, a_d + a_{d+1} + \cdots + a_m, 0, \ldots, 0).$$

It is easily verified that changing one coordinate in an m-tuple changes one coordinate in its image, so f_d is an endomorphism; its image is the Hamming graph $H(d,k)$ with k^d vertices, as required.

So for this group, there are maps which are not synchronized but have a variety of ranks, not just the minimum possible rank. However, all these maps have uniform kernels. Can there exist a map with non-uniform kernel synchronized by G?

These considerations suggested the following definition. Let us call a transitive permutation group G on Ω *almost synchronizing* if every map on Ω which is not uniform (i.e. not all its kernel classes have the same size) is synchronized by G. We note that an almost synchronizing group is primitive. For suppose that G is imprimitive, and preserves the non-trivial partition P. Form the multipartite graph in which two vertices are adjacent if they lie in different parts of P. This graph is preserved by G, but we can collapse vertices within each part of P arbitrarily by endomorphisms.

On the strength of this observation and examples like those above, it was conjectured for some time that every primitive group is almost synchronizing.

This was very recently shown to be false [11]. We describe here a general construction from that paper.

Consider the following two graph products. Let Γ and Δ be graphs with vertex sets A and B respectively.

(a) The *cartesian product* $\Gamma \,\square\, \Delta$ has vertex set $A \times B$; there is an edge from (a_1, b_1) to (a_2, b_2) if either $a_1 = a_2$ and b_1 is adjacent to b_2 in Δ, or a_1 is adjacent to a_2 in Δ and $b_1 = b_2$.

(b) The *categorical product* $\Gamma \times \Delta$ also has vertex set $A \times B$; but there is an edge from (a_1, b_1) to (a_2, b_2) if there are edges from a_1 to a_2 in Γ and from b_1 to b_2 in Δ.

The notation for these products is chosen so that the product symbol represents the corresponding product of two edges.

For example, the Cartesian product of two copies of K_r is the Hamming graph $H(2, r)$, while the categorical product is the complement of $H(2, r)$.

Now here is a flexible construction of primitive graphs with non-uniform endomorphisms.

Example 5.1. Let Γ be a graph whose automorphism group acts primitively on its vertices. Then the Cartesian product $\Gamma \,\square\, \Gamma$ is also vertex-primitive, with automorphism group $\mathrm{Aut}(\Gamma) \mathrm{wr} S_2$. In addition, if the chromatic and clique number of Γ are both equal to k, then $V(\Gamma)$ can be partitioned into k colour classes, say V_1, V_2, ..., V_k; these classes have equal size, by Theorem 3.9. So there is a surjective homomorphism $\Gamma \,\square\, \Gamma \to K_k \,\square\, K_k$ with kernel classes $V_i \times V_j$ for $1 \leq i, j \leq k\}$. Moreover, there is a homomorphism from Γ to $\Gamma \,\square\, \Gamma$; simply take the second coordinate to be fixed.

If in addition there is a homomorphism $f : K_k \,\square\, K_k \to \Gamma$, then by composing homomorphisms

$$\Gamma \,\square\, \Gamma \to K_k \,\square\, K_k \xrightarrow{f} \Gamma \to \Gamma \,\square\, \Gamma,$$

there is an endomorphism of $\Gamma \,\square\, \Gamma$. Moreover, if the homomorphism f is non-uniform, then the endomorphism is also non-uniform; and its rank is equal to the rank of f.

We can obtain examples by taking Γ to be the complement of $K_k \,\square\, K_k$, that is, $\Gamma = K_k \times K_k$. Now a homomorphism f from $K_k \,\square\, K_k$ to $K_k \times K_k$ is given by $(u, v) \mapsto (g(u, v), h(u, v))$, where the two coordinate functions $g(u, v)$ and $h(u, v)$ satisfy the homomorphism requirement that if (u, v) and (u', v') agree in one position but not the other, then $g(u, v) \neq g(u', v')$ and $h(u, v) \neq h(u', v')$.

In other words, g and h are *Latin squares* of order k: that is, they define $k \times k$ arrays with entries from a set of size k such that no entry is repeated in a row or a column. But note that there is no connection between the two Latin squares!

The rank of the homomorphism is the number of ordered pairs of symbols which arise when the two Latin squares are superimposed. The possibilities have been determined by Colbourn, Zhu and Zhang [43, 116]:

Theorem 5.2. *There are two Latin squares of order k whose superposition gives r ordered pairs of symbols if and only if $r = k$, or $r = k^2$, or $k + 2 \leq r \leq k^2 - 2$, with the following exceptions:*

(a) $k = 2$ *and* $r = 4$*;*

(b) $k = 3$ *and* $r \in \{5, 6, 7\}$*;*

(c) $k = 4$ *and* $r \in \{7, 10, 11, 13, 14\}$*;*

(d) $k = 5$ *and* $r \in \{8, 9, 20, 22, 23\}$*;*

(e) $k = 6$ *and* $r \in \{33, 36\}$.

Note that the case $r = k$ corresponds to using the same Latin square twice, while $r = k^2$ corresponds to a pair of orthogonal Latin squares. In these cases, the endomorphism constructed is uniform; but in general it is not.

We conclude that, for each value of r with $r = k$, or $r = k^2$, or $k+2 \leq r \leq k^2-2$, with the exceptions given in Theorem 5.2, there is a map of rank k not synchronized by the primitive group $G = (S_k \text{ wr } C_2) \text{ wr } C_2$ (of degree k^4); in most cases, these maps are not uniform, so the group G is not almost synchronizing.

Here is another example from [11]. Though this construction is not as flexible as the previous one, it was the first one found, produces a primitive group of smallest possible degree (namely 45) which is not almost synchronizing, and also produces a non-uniform map of smallest possible rank (namely 5) which fails to be synchronized by a primitive group.

We start with three particular graphs. Two of these examples are two of the three "remarkable graphs" discussed by Biggs [25], namely the Petersen graph (see Figure 1) and the Biggs–Smith graph; the third is the *Tutte–Coxeter graph* on 30 vertices [107, 44]. All are trivalent graphs without triangles, and have proper 3-edge colourings; and their automorphism groups act primitively on the edges. We consider their line graphs. These are 4-valent vertex-primitive graphs with chromatic number 3; the closed neighbourhood of a vertex is the *butterfly graph* shown in Figure 8.

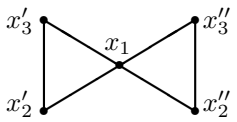

Figure 8. The butterfly

Take a 3-colouring of one of these line graphs, with colour classes C_1, C_2, C_3. A colour class contains no edges; so if v is a vertex in C_3, then two neighbours of v lie in each of the other two classes. So the induced subgraph on $C_2 \cup C_3$ has valency 2, and is a union of cycles. (The cycles have even length since the colours alternate.) If this induced subgraph is disconnected, let C' and C'' be unions of connected components; and let $C_2' = C' \cap C_2$, $C_2'' = C'' \cap C_2$, and similarly for C_3' and C_3''. This gives partitions $C_2 = C_2' \cup C_2''$ and $C_3 = C_3' \cup C_3''$, with edges only between C_2' and C_3' and between C_2'' and C_3''. In other words, we have a homomorphism from $L(\Gamma)$ onto the butterfly, where C_1 maps to the "body" x, and C_i' and C_i'' to x_i' and x_i'' respectively for $i = 2, 3$ (see Figure 8).

The line graph of the Petersen graph has a unique 3-colouring up to isomorphism, and $C_2 \cup C_3$ turns out to be connected. But in each of the other cases, the required homomorphism exists. For the line graph of the Tutte–Coxeter graph (with 45 vertices), the induced subgraph on $C_2 \cup C_3$ falls into components of sizes 10 and 20, and so we obtain a homomorphism with kernel classes of sizes 5, 5, 10, 10, and 15. Since the butterfly is a subgraph, we can realise this map as an endomorphism of $L(\Gamma)$; so the automorphism group of $L(\Gamma)$ is primitive but not almost synchronizing.

Further analysis shows that this graph has an endomorphism onto the "double butterfly" (a triangle with triangles attached at two of its vertices) with rank 7. (We performed the computation using the software systems GAP [56] and MINION [55].)

The line graph of the Biggs–Smith graph is even more prolific. It has non-uniform endomorphisms of ranks 5, 7 and 9 (the last onto the "triple butterfly", a triangle with triangles attached at each vertex).

An almost synchronizing group is primitive (since an imprimitive group preserves a complete multipartite graph, which has non-uniform endomorphisms). However, it need not be basic: the automorphism group of the square grid (Figure 6) fails to synchronize only uniform maps of rank 3, but is clearly not basic. In the other direction, the butterfly examples in this section show that primitive groups which fail to be almost synchronizing may or may not be basic. So there is no implication between these concepts.

5.2. Separating groups. The next variant is based on the following theorem. A set of vertices of a graph Γ is *independent* if it contains no edge (so the induced subgraph on the set is a null graph). Let $\alpha(\Gamma)$ be the *independence number* of Γ, the size of the largest independent set. Note that $\alpha(\Gamma) = \omega(\overline{\Gamma})$.

Theorem 5.3. *Let Γ be a graph on n vertices, whose automorphism group G acts transitively on vertices. Then $\alpha(\Gamma) \cdot \omega(\Gamma) \leq n$. Equality holds if and only if any clique and any independent set of maximum size have non-empty intersection.*

Proof. Let A be a clique and B an independent set. Count triples a, b, g with $a \in A$, $b \in B$, and $ag = b$. There are $|A|$ choices for a, $|B|$ choices for B, and $|G|/|\Omega|$ choices for g (since the set of such g is a coset of the stabilizer of a point, and all such cosets have the same cardinality). On the other hand, there are $|G|$ choices for g, and for each choice there is at most one element in $Ag \cap B$, since a clique and an independent set cannot have more than one point in common. So

$$|A| \cdot |B| \cdot |G|/n \leq |G|,$$

giving $|A| \cdot |B| \leq n$. If equality holds, then $Ag \cap B \neq \emptyset$ for all $g \in G$, and in particular $A \cap B \neq \emptyset$; this holds for any clique and independent set. □

Let G be a transitive permutation group on Ω. If A and B are subsets of Ω which satisfy $|A| \cdot |B| < |\Omega|$, then A and B can be *separated* by G: that is, there

exists $g \in G$ such that $Ag \cap B = \emptyset$. (This was proved by Peter Neumann [83]; an elementary proof appears in [26].) What happens if $|A| \cdot |B| = |\Omega|$?

We say that the transitive group G on Ω is *non-separating* if there exist two subsets A, B of Ω (other than singleton sets and Ω) satisfying $|A| \cdot |B| = |\Omega|$, for which $|Ag \cap B| = 1$ for all $g \in G$. If this is not the case, then since the average size of this intersection is 1 (see Theorem 5.12 below for a proof of a generalization of this), there must be some $g \in G$ for which $Ag \cap B = \emptyset$. We say that G is *separating* if, for any choice of two sets A and B with $|A| \cdot |B| = |\Omega|$ and $|A|, |B| > 1$, there exists $g \in G$ with $Ag \cap B = \emptyset$.

Theorem 5.4. *Let G be a transitive permutation group on Ω. Then G is non-separating if and only if there is a non-trivial graph Γ on the vertex set Ω satisfying $\omega(\Gamma)\alpha(\Gamma) = |\Omega|$ and $G \leq \mathrm{Aut}(\Gamma)$.*

Proof. If a graph Γ with these properties exists, then taking A and B to be a clique and an independent set of maximum size, we see that G cannot separate them; so it is not separating.

Conversely, suppose that G is not separating, and let A and B be the sets specified in the definition, so that $|A| \cdot |B| = |\Omega|$. Now let Γ be the graph on Ω whose edges are all the images under G of pairs of vertices in A. Then $G \leq \mathrm{Aut}(\Gamma)$; A is a clique in Γ; and, since no element of G maps A to a set intersecting B in two points, B is an independent set in Γ. □

Corollary 5.5. *A separating group is synchronizing.*

Proof. Suppose that G is not synchronizing, and let Γ be a graph with clique number equal to chromatic number which witnesses this (see Theorem 4.4). Then we have $G \leq \mathrm{Aut}(\Gamma)$. Each colour class is an independent set, and the average size of the colour classes is $|\Omega|/|\chi(\Gamma)| = |\Omega|/|\omega(\Gamma)|$; but no independent set can exceed this size, so $\alpha(\Gamma) = |\Omega|/|\omega(\Gamma)|$.

Alternatively, if P is a section-regular partition with section A and B is any part of P, then $|A| \cdot |B| = |\Omega|$ and $|Ag \cap B| = 1$ for all $g \in G$ (see Theorem 3.8). □

The converse of this corollary is false; we will see an example later.

A transitive group of prime degree is obviously separating since there are no non-trivial subsets A, B with $|A| \cdot |B| = |\Omega|$.

Corollary 5.6. *A transitive group of prime degree is separating and hence synchronizing.*

We can now make a small improvement in the algorithm for testing synchronization. As a comment on this, both clique number and chromatic number are NP-hard, but in practice finding clique number is very much easier than finding chromatic number. (There are results in the theory of parameterized complexity which support this assertion. The k-clique problem lies in the complexity class W[1] and is complete for this class – see Downey and Fellows [48] – but k-colouring is NP-complete even for $k = 3$.)

We modify Algorithm 4.6 as follows.

Algorithm 5.7. (a) The $2^r - 2$ non-trivial G-invariant graphs (where r is the number of G-orbits on 2-element subsets of Ω) fall into $2^{r-1} - 1$ complementary pairs.

(b) For each pair, find the clique numbers of the two graphs in the pair. If their product is $|\Omega|$, then G is not separating; remember this pair of graphs. If this never happens, then G is separating (and hence synchronizing).

(c) Now we just have to look at the graphs produced in the second stage of the algorithm, and check whether they have clique number equal to chromatic number (noting that we now know the clique number). If this never happens, then G is synchronizing; otherwise, not.

Example 5.8. Consider the symmetric group S_m acting on 2-sets. We have just one complementary pair of graphs to consider: the line graph of K_m (which has clique number $m - 1$) and its complement (which has clique number $\lfloor m/2 \rfloor$). We see immediately that this group is separating if (and only if) m is odd. In the case m even, we have to work out the chromatic number of these two graphs, as we did earlier, and we find that this group is not synchronizing.

So, for these groups, the properties "synchronizing" and "separating" are equivalent.

5.3. Partition separation. We have seen that separation is a strengthening of synchronization. There is a dual notion, which is a weakening of synchronization. We say that a transitive permutation group G on Ω is *not partition-separating* if there are two non-trivial partitions P and Q of Ω such that, for any part A of P and any part B of Q, and any $g \in G$, $|Ag \cap B| = 1$. (This implies that each of P and Q is section regular, and any part of one is a G-section for the other.) It is *partition-separating* if no such pair of partitions exists. Note that, if P and Q witness the failure of partition-separation, then the partitions P and Q are uniform, and the number of parts of P is equal to the size of a part of Q and *vice versa*.

By very similar arguments to those we have seen, we obtain the following (previously unpublished) characterization:

Theorem 5.9. *The transitive permutation group G on Ω is not partition-separating if and only if there exists a non-trivial graph Γ on the vertex set Ω with $G \leq \mathrm{Aut}(\Gamma)$ and $\chi(\Gamma) \cdot \chi(\overline{\Gamma}) = |\Omega|$.*

A transitive group G which is partition-separating is primitive. For if G is imprimitive, the above theorem applies to the complete multipartite graph whose parts are those of the non-trivial partition fixed by G. Furthermore, we have:

Proposition 5.10. *A partition-separating group is basic.*

Proof. Consider the Cartesian structure with automorphism group $S_k \,\mathrm{wr}\, S_m$, and identify the domain of S_k with the group of integers mod k. Now consider the following two partitions:

- The parts of the first partition consist of all m-tuples where the values of all coordinates except the last are constant: there are k^{m-1} parts of size k.
- The parts of the second partition are the m-tuples with fixed sum: there are k parts of size k^{m-1}.

It is straightforward to show that any part of one is a section for the other. □

Example 5.11. We have seen that the symmetric group S_m acting on 2-sets is not synchronizing if m is even and greater than 4. However, this group is partition-separating. For the line graph of K_n has clique number equal to chromatic number in this case, but its complement does not.

5.4. Multisets. In order to describe the next class of permutation groups, the *spreading* groups, we need to introduce some notation for multisets.

A *multiset* of Ω is a function from Ω to the natural numbers (including zero). If A is a multiset, we call $A(i)$ the *multiplicity* of i in A. The set of elements of Ω with non-zero multiplicity is the *support* of A. By abuse of language, we say that i belongs to A if it belongs to the support of A. We can regard a set as a special multiset in which all multiplicities are zero and one (identifying the set with its characteristic function).

The *cardinality* of A is

$$|A| = \sum_{i\in\Omega} A(i);$$

this agrees with the usual definition in the case of a set.

The *product* of two multisets A and B of Ω is the multiset $A * B$ defined by

$$(A * B)(i) = A(i)B(i).$$

This is a generalization of the usual definition of intersection of sets; but the "intersection" of multisets is defined differently in the literature.

(a) The product of two sets is their intersection.

(b) The product of a multiset A and a set B is the "restriction of A to B", that is, points of B have the same multiplicity as in A, while points outside B have multiplicity zero.

(c) if we identify a multiset A with a vector v_A of non-negative integers with coordinates indexed by Ω, then we have $|A * B| = v_A \cdot v_B$ for all multisets A and B. In particular, $|A| = v_A \cdot j$, where j is the all-one vector.

The image of a multiset A under a permutation g is defined by

$$Ag(i) = A(ig^{-1}).$$

This agrees with the usual image of a set under a permutation. We say that $a \in A$ if the multiplicity of $A(a)$ is at least 1.

Theorem 5.12. *Let G be a transitive permutation group on Ω, and let A and B be multisets of Ω. Then the average cardinality of the product of A and Bg is given by*

$$\frac{1}{|G|}\sum_{g\in G}|A * Bg| = \frac{|A|\cdot|B|}{|\Omega|}.$$

Proof. We count triples (a, g, b) with $a \in A$, $g \in G$, $b \in B$, and $bg = a$. (Points of A or B are counted according to their multiplicity.) There are $|A|$ choices for a and $|B|$ choices for b. Then the set of elements of G mapping b to a is a right coset of the stabilizer G_b since G is transitive, so there are $|G|/|\Omega|$ such elements.

On the other hand, for each element $g \in G$, if $bg = a$, then this element belongs to the support of $A * Bg$. The number of choices of a is equal to the sum of multiplicities in A, and for each one, the number of choices of b is the multiplicity of ag^{-1} in B, that is, of a in Bg. So the product counts the multiplicities correctly.

Equating the two sides gives the result. □

5.5. Spreading. The concept of spreading has been used by various authors in studying the Černý conjecture (for example, [14, 86]) but the general definition appears in print here for the first time.

Let G be a transitive permutation group on Ω, and A and B multisets of Ω. Consider the following four conditions, where λ is a positive integer:

$(1)_\lambda$: $|A * Bg| = \lambda$ for all $g \in G$.

(2): A is a set.

(3): B is a set.

(4): $|A|$ divides $|\Omega|$.

Note that

(a) $(1)_\lambda$ is symmetric in A and B.

(b) $(1)_\lambda$ with $\lambda = 1$ implies (2), (3) and (4). For, if $A(i) > 1$, then choosing g to map a point in the support of B to i, we would have $|A \cap Bg| > 1$; so (2) holds, and (3) is similar. Finally, if $(1)_\lambda$ holds with $\lambda = 1$ then $|A|\cdot|B| = |\Omega|$ by Theorem 5.12.

(c) If (2) and (3) hold, then we can replace product by intersection in $(1)_\lambda$.

We will call a multiset *trivial* if either it is constant or its support is a singleton. (This is a slight departure from our previous convention on non-triviality!)

The transitive permutation group G on Ω is *non-spreading* if there exist non-trivial multisets A and B and a positive integer λ such that $(1)_\lambda$, (3) and (4) hold, and is *spreading* otherwise. Note that if $(1)_\lambda$ holds, then

$$\lambda = \frac{|A|\cdot|B|}{|\Omega|} \tag{1}$$

by Theorem 5.12.

Theorem 5.13. *The permutation group G on Ω is spreading if and only if, for any function $t\colon \Omega \to \Omega$ which is not a permutation and any non-trivial subset S of Ω, there exists $g \in G$ such that $|Sgt^{-1}| > |S|$.*

Proof. Suppose that G is non-spreading, and let the multiset A and set B be witnesses. Since $|A|$ divides $|\Omega|$, there is a function t from Ω to Ω so that $|at^{-1}|$ is proportional to the multiplicity of a in A (the constant of proportionality being $|\Omega|/|A|$). Let $S = B$. Then for any $g \in G$, we have

$$|Sgt^{-1}| = |A * Sg| \cdot |\Omega|/|A| = |S|,$$

by the definition of non-spreading and equation (1).

Conversely, suppose that there is a function t and subset S for which the condition in the theorem is false. Let A be the multiset in which the multiplicity of a is equal to $|at^{-1}|$. Then we have $|A| = |\Omega|$ and it is false that $|A * Sg| > |S|$ for any $g \in G$; thus we have $|A * Sg| = |S|$ for all $g \in G$ (since the average value of $|A * Sg|$ is $|S|$ by Theorem 5.12). We conclude that $(1)_{|S|}$, (3) and (4) hold, so that G is non-spreading. □

Theorem 5.14. (a) *A spreading permutation group is separating.*

(b) *A 2-homogeneous group is spreading.*

Proof. (a) Witnesses to non-separation are also witnesses to non-spreading (with $\lambda = 1$).

(b) The arguments are similar to those we have seen before. □

We will see that neither implication reverses. In fact, in our terminology, Pin proved that transitive groups of prime degree are spreading [86]. We shall obtain this as a special case of a stronger result later.

Example 5.15. We saw that the automorphism group of the Petersen graph is synchronizing. However, this group is not spreading. Take A to be the outer pentagon, and B an independent set of size 4: then $|Ag \cap B| = 2$ for any automorphism g.

More generally, we saw in Example 5.8 that S_n, acting on the set of 2-subsets of $\{1, \ldots, n\}$, is separating if n is odd and $n \geq 5$. We now show that it is not spreading if n is odd.

Let A be a set of n pairs forming an n-cycle: $A = \{\{1,2\}, \{2,3\}, \ldots, \{n-1, n\}, \{n, 1\}\}$. Let B be the set of $n-1$ pairs containing the fixed element 1. Then

(a) $|Ag \cap B| = 2$ for all $g \in G$;

(b) A and B are sets;

(c) $|A| = n$ divides $|\Omega| = n(n-1)/2$ if n is odd.

5.6. Spreading groups and the Černý conjecture.

Theorem 5.16. *Let G be a spreading permutation group on Ω, and f a function from Ω to Ω which is not a permutation. Then $\langle G, f\rangle$ contains a rank 1 mapping that can be expressed as a product which has at most $n-1$ occurrences of f, where $n = |\Omega|$.*

In particular, if A is a generating set for G and each element of G can be written as a word in A of length at most t, then there is a reset word over $A \cup \{f\}$ of length at most $1 + (t+1)(n-2)$.

In other words, the property of being spreading not only implies synchronization, but also realizes the first part of our programme for bounding the length of the reset word.

Proof. Suppose that we have a set U_k with $|U_k| \geq k$, such that there is a word w in $\langle G, f\rangle$ with at most $k-1$ occurrences of f which maps U_k to a singleton.

By Theorem 5.13, there exists $g \in G$ such that $U_{k+1} = U_k g f^{-1}$ satisfies $|U_{k+1}| \geq k+1$. We have $U_k = U_{k+1} f g^{-1}$, so the word $fg^{-1}w$ with at most k occurrences of f maps U_{k+1} to a singleton.

By induction on k, the result is proved.

The final statement follows because there is a rank 1 mapping of the form $fg_1fg_2\cdots fg_{n-2}f$ and each mapping fg_i can be represented by a word of length at most $t+1$. □

5.7. Measuring non-synchronization. Given a permutation group G on Ω, we want to give a quantitative measure of how far (if at all) G is from being synchronizing.

There are several methods for doing this. For example, we could consider the largest and smallest rank of a map not synchronized by G. Theorem 2.8 shows that G is primitive if and only if it synchronizes every map of rank $n-1$ (where $n = |\Omega|$). As we noted after the proof of the version of this result included in Theorem 4.8, it was shown in [11] that primitive groups synchronize maps of rank at least $n-4$. At the other extreme, we have the following, due to Peter Neumann [84]:

Theorem 5.17. *A primitive group synchronizes every map of rank 2.*

Proof. Suppose that G fails to synchronize a map f of rank 2. Then 2 is the minimum rank of a map in $\langle G, f\rangle$. By Theorem 4.4, there is a graph Γ with clique number and chromatic number 2 (that is, a non-trivial bipartite graph) with $G \leq \mathrm{Aut}(\Gamma)$. If Γ is disconnected, then the partition into connected components is preserved by G; if it is connected, then it has a unique bipartition, which is preserved by G. □

The example of the 3×3 grid (Figure 6) shows that this result does not extend to maps of rank 3.

More generally, given a group G, we define the set

$$\mathrm{NS}(G) = \{r : \text{there exists a map of rank } r \text{ not synchronized by } G\}$$

of *non-synchronizing ranks* for G.

Theorem 5.18. *If G is transitive but imprimitive, of degree n, then*

$$|\,\mathrm{NS}(G)| \geq (\tfrac{3}{4} + o(1))n.$$

Proof. Suppose that G has m blocks of imprimitivity, each of size k. Then, among the non-trivial G-invariant graphs, we find:

(a) The disjoint union of m complete graphs of size k. This graph can be mapped onto any non-empty subset of its components; so

$$\{k, 2k, \ldots, (m-1)k\} \subseteq \mathrm{NS}(G).$$

(b) The complete multipartite graph with m parts of size k. Each part can be collapsed onto any non-empty subset of itself; so

$$\{m, m+1, m+2, \ldots, mk-2, mk-1\} \subseteq \mathrm{NS}(G).$$

It is easy to see that the union of these two subsets has size $(\frac{3}{4} - o(1))n$. □

Conjecture 5.19. If G is primitive of degree n, then $|\,\mathrm{NS}(G)| = o(n)$.

If true, this would show that, as far as synchronization is concerned, there is a big divide between primitive and imprimitive groups, with primitive groups being close to synchronizing, and imprimitive groups more distant. The most extreme primitive groups known, constructed in Example 5.1, have $|\,\mathrm{NS}(G)| = O(\sqrt{n})$ ([11]).

6. Examples

In this section, we treat some general classes of examples. These will yield examples of groups which are synchronizing but not separating. We will see that

(a) the techniques are combinatorial and geometric rather than group-theoretic;

(b) we reach very hard problems very quickly.

6.1. The symmetric group on subsets. Let $G = S_n$, and let Ω be the set of all k-subsets of $\{1, \ldots, n\}$.

We may assume that $n \geq 2k$, since the actions of S_n on k-sets and on $(n-k)$-sets are isomorphic.

In fact we may assume that $n \geq 2k+1$, since the action of S_n on k-sets is imprimitive if $n = 2k$: the relation "equal or disjoint" is a congruence.

Now G has k orbits on the 2-element subsets of Ω, namely,

$$O_l = \{\{S_1, S_2\} : |S_1 \cap S_2| = l\}$$

for $l = 0, 1, \ldots, k-1$. These k graphs together with the relation of equality form a combinatorial structure known as an *association scheme*, specifically the *Johnson scheme* $J(n,k)$. (Association schemes will be discussed further in Section 10.1.)

All these graphs are connected (this is an exercise), so G is primitive on Ω. Since its socle is simple, it is basic.

If $k = 1$, then G is 2-transitive. We ignore this case. Also, we dealt with the case $k = 2$ earlier. So we assume that $k \geq 3$.

6.1.1. Baranyai's Theorem. Let $\mathcal{F}$ be a set of k-subsets of $\{1, \ldots, n\}$, where k divides n. A 1-*factorization* of $\mathcal{F}$ is a partition of $\mathcal{F}$ such that each part is a partition of $\{1, \ldots, n\}$ (that is, a set of n/k pairwise disjoint subsets).

Theorem 6.1. *If k divides n, then there is a 1-factorization of the set of all k-subsets of $\{1, \ldots, n\}$.*

The theorem was proved by Baranyai in 1973 ([19]). The proof is a beautiful application of the Max-Flow Min-Cut Theorem for networks.

As a corollary we have:

Theorem 6.2. *If k divides n, then S_n acting on k-sets is not synchronizing.*

For the set of all k-sets containing a fixed element (say 1) is a section of the Baranyai partition, which is thus section-regular.

6.1.2. The case $k = 3$. We now consider the case $k = 3$, and resolve completely the question of synchronization and separation. We will see that further combinatorial tools are required.

Theorem 6.3. *Let $G = S_n$ acting on the set of 3-subsets of $\{1, \ldots, n\}$, with $n \geq 7$. Then the following are equivalent:*

(a) *G is synchronizing;*

(b) *G is separating;*

(c) *n is congruent to 2, 4 or 5 (mod 6), and $n \neq 8$.*

Note that synchronization and separation are equivalent for this class of groups.

A *Steiner triple system* is a collection $\mathcal{S}$ of 3-subsets of $\{1, \ldots, n\}$ with the property that every pair of points of $\{1, \ldots, n\}$ is contained in a unique member of $\mathcal{S}$.

Kirkman proved in 1847 that a Steiner triple system on n points exists if and only if n is congruent to 1 or 3 mod 6.

A *large set* of Steiner triple systems is a partition of the set of all 3-subsets of $\{1, \ldots, n\}$ into Steiner triple systems. (Counting shows that there must be $n-2$ such systems.)

For $n = 7$, there is a unique Steiner triple system, the *Fano plane* (see Figure 9.)

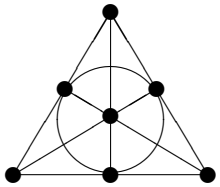

Figure 9. The Fano plane

We cannot find more than two disjoint copies of the Fano plane. This fact goes back to Cayley. However, Teirlinck [103] showed:

Theorem 6.4. *If n is congruent to 1 or 3 (mod 6) and $n > 7$, then there exists a large set of Steiner triple systems on n points.*

Now let G be S_n acting on 3-sets, for $n \geq 7$.

Baranyai's theorem shows that G is non-synchronizing if n is divisible by 3, that is, if n is congruent to 0 or 3 (mod 6).

Teirlinck's theorem shows that G is non-synchronizing if n is congruent to 1 or 3 (mod 6) and $n \neq 7$. (The set of triples through two given points is a section for all images of the large set.)

The cases $n = 7$ and $n = 8$ require special treatment.

The case $n = 7$ For each line L of the Fano plane, let $S(L)$ be the set of 3-sets equal to or disjoint from L. Then $|S(L)| = 5$.

Since no two lines of the Fano plane are disjoint, and no 3-set is disjoint from more than one line, we see that the sets $S(L)$ are pairwise disjoint. Since $5 \cdot 7 = 35 = \binom{7}{3}$, they form a partition of Ω.

Now 3-sets in the same $S(L)$ meet in 0 or 2 points. So any image of the Fano plane meets each $S(L)$ in at most (and hence exactly) one set. Thus the partition is section-regular, the Fano plane being the section.

So S_7 acting on 3-sets is not synchronizing.

The case $n = 8$ Take a Fano plane on $\{1, \ldots, 7\}$. For each line L of the Fano plane, partition the eight points into $L \cup \{8\}$ and the rest, and take the set $T(L)$ of eight triples contained in a part of this partition. This gives a partition of all the $\binom{8}{3} = 56 = 7 \cdot 8$ 3-sets into seven subsets of size 8.

Once again we find that this partition is section-regular, with the Fano plane as a section.

6.1.3. The separating cases. We have now shown that, in the cases not stated in the theorem, G is non-synchronizing and hence non-separating. We have to show that, in the remaining cases, G is separating, and hence synchronizing.

There are $2^3 - 2$ graphs to consider. We denote them by Γ_I, for $\emptyset \subset I \subset \{0, 1, 2\}$; the vertices are the 3-sets, and two vertices are adjacent if and only if the cardinality of their intersection belongs to I.

According to Theorem 5.4, we have to find the clique number of each of these graphs, and check whether $\omega(\Gamma_I)\omega(\Gamma_{I^*}) = \binom{n}{3}$, where $I^* = \{0, 1, 2\} \setminus I$.

The following theorem, the *Erdős–Ko–Rado theorem*, finds the clique number of some of these graphs. A family $\mathcal{F}$ of k-subsets of $\{1, \ldots, n\}$ is *t-intersecting* if $|A \cap B| \geq t$ for all $A, B \in \mathcal{F}$.

Theorem 6.5. *For $n \geq n_0(k, t)$, the maximum size of a t-intersecting family of k-sets of $\{1, \ldots, n\}$ is $\binom{n-t}{k-t}$, with equality realized only by the family of all k-sets containing a fixed t-set.*

The correct value of $n_0(k, t)$ is known. We need only that the assertion of the theorem is true for $k = 3$, $n \geq 7$, and $t = 1$ or $t = 2$.

The cases $I = \{0\}$ and $I = \{1, 2\}$ Clearly $\omega(\Gamma_{\{0\}}) = \lfloor n/3 \rfloor$.

By Erdős–Ko–Rado, $\omega(\Gamma_{\{1,2\}}) = \binom{n-1}{2}$. The product of these numbers is $\binom{n}{3}$ if and only if n is a multiple of 3; but this case is excluded.

The cases $I = \{0, 1\}$ and $I = \{2\}$ By Erdős–Ko–Rado, $\omega(\Gamma_{\{2\}}) = n - 2$.

A clique in $\Gamma_{\{0,1\}}$ has the property that two points lie in at most one set in the clique; so $\omega(\Gamma_{\{0,1\}}) \leq n(n-1)/6$, with equality if and only if there is a Steiner triple system of order n, that is, n is congruent to 1 or 3 (mod 6). But these cases are excluded.

The cases $I = \{1\}$ and $I = \{0, 2\}$ It is easy to show that a maximum clique in $\Gamma_{\{0,2\}}$ is obtained by dividing most of $\{1, \ldots, n\}$ into disjoint 4-sets and taking all the 3-subsets of these 4-sets. In particular, $\omega(\Gamma_{\{0,2\}}) \leq n$.

A maximum clique in $\Gamma_{\{1\}}$ is obtained by taking $\lfloor n/2 \rfloor$ triples through a fixed point but having no further point in common, provided that $n \geq 17$. For smaller values, a Fano plane may be better.

A little calculation shows that the product of these bounds is strictly smaller than $\binom{n}{3}$ except for $n = 7$ and $n = 8$; but these cases are excluded.

6.1.4. Spreading. The recent remarkable result of Keevash [72] on the existence of Steiner systems shows, as above, the existence of infinitely many more values of n and k for which the symmetric group S_n acting on k-sets is non-separating.

However, for spreading, things are much easier. The following argument is due to Peter Neumann.

Theorem 6.6. *The symmetric group S_n acting on k-sets is always non-spreading.*

Proof. Let d be the greatest common divisor of n and k. Let H be a cyclic group of order n permuting the elements of $\{1,\dots,n\}$ in the natural way. Now choose a k-subset of $\{1,\dots,n\}$ which is a union of k/d orbits of the subgroup of order d of H, and let A be the H-orbit (in Ω) containing this set; so $|A| = n/d$. Let B consist of all k-sets containing the element 1. Since A is invariant under a transitive group, $|A \cap Bg|$ is constant for $g \in G$. Also, clearly A and B are sets.

It remains only to show that $|A| = n/d$ divides $|\Omega| = \binom{n}{k}$. The stabilizer in H of any k-set has order dividing k and also dividing n, hence dividing d; so the size of any H-orbit in Ω is a multiple of n/d. The assertion follows. □

6.1.5. Linear groups acting on subspaces. The action of $\mathrm{PGL}(n,q)$ on the set of k-dimensional subspaces of the n-dimensional vector space gives a linear analogue of the action of S_n on k-subsets of $\{1,\dots,n\}$. But much less is known in this case, since the linear analogues of combinatorial results such as those of Baranyai and Erdős–Ko–Rado are not known except in special cases (for example, [24]). Even the existence of the analogues of Steiner systems is a major unsolved problem; the first examples have been given very recently [30].

6.2. Classical groups and polar spaces. Now we turn to the other family of examples discussed here: classical (symplectic, unitary and orthogonal) groups acting on the associated polar spaces.

We give a brief introduction to these groups and geometries; more detail is available in several places, including [34, 102].

We are only interested in finite classical groups; this makes the theory simpler in several respects.

6.2.1. Finite classical groups. A classical group acts on a vector space and preserves a form of some type:

(a) for *symplectic groups*, an alternating bilinear form;

(b) for *unitary groups*, a Hermitian sesquilinear form;

(c) for *orthogonal groups*, a quadratic form, and the symmetric bilinear form obtained from it by *polarization*.

The basic form should be *non-degenerate* or *non-singular*. The reason for separating cases is that strange things happen with quadratic forms in characteristic 2. But we can ignore this complication!

There are three parameters associated with a classical group:

q, the order of the field over which the matrices are defined;

r, the *Witt index*, the dimension of the largest subspace on which the form vanishes identically;

ϵ, a parameter defined shortly.

We denote the dimension of the underlying vector space by n.

We divide the classical groups into six families:

symplectic: $\mathrm{PSp}(2r, q)$, $n = 2r$

unitary: $\mathrm{PSU}(2r, q_0)$, $n = 2r$, and $PSU(2r + 1, q_0)$, $n = 2r + 1$;

orthogonal: $\mathrm{P}\Omega^+(2r, q)$, $n = 2r$; $\mathrm{P}\Omega(2r+1, q)$, $n = 2r+1$; and $\mathrm{P}\Omega^-(2r+2, q)$, $n = 2r + 2$.

Note that for the unitary groups, the field order must be a square, say $q = q_0^2$, and there is a field automorphism $x \mapsto x^{q_0}$ of order 2. We use the group-theorists' notation $\mathrm{PSU}(n, q_0)$, but the field of definition is $\mathbb{F}_q$.

We need not consider orthogonal groups of odd dimension over fields of characteristic 2, since they turn out to be isomorphic to symplectic groups of one dimension less.

The values of the parameter ϵ are given in the table:

Type	ϵ
$\mathrm{PSp}(2r, q)$	0
$\mathrm{PSU}(2r, q_0)$	$-\frac{1}{2}$
$\mathrm{PSU}(2r + 1, q_0)$	$\frac{1}{2}$
$\mathrm{P}\Omega^+(2r, q)$	-1
$\mathrm{P}\Omega(2r + 1, q)$	0
$\mathrm{P}\Omega^-(2r + 2, q)$	1

6.2.2. Polar spaces. The *polar space* associated with a classical group acting on a vector space V is the geometry of *totally isotropic* subspaces of V, those on which the form vanishes identically. We abbreviate this to *t.i.*

In the case of orthogonal groups, we should really use the term *totally singular* or *t.s.* instead; but we will ignore this distinction.

Subspaces of (vector space) dimension 1 or 2 are called *points* and *lines*, as usual in projective geometry. Subspaces of maximum dimension r are called *maximal subspaces*.

6.2.3. Numerical information. Numerical information about polar spaces can be expressed in terms of the parameters q, r, ϵ:

Theorem 6.7. (a) *The number of points of the polar space is* $(q^r - 1)(q^{r+\epsilon} + 1)/(q - 1)$*; each maximal subspace contains* $(q^r - 1)/(q - 1)$ *points.*

(b) *The number of points not collinear with a given point is* $q^{2r+\epsilon-1}$.

(c) *The number of maximal subspaces is*

$$\prod_{i=1}^{r}(1+q^{i+\epsilon}).$$

6.2.4. Witt's Lemma. *Witt's Lemma* asserts that the action of the classical group on a polar space is "homogeneous", in the sense that any linear isometry between subspaces of the vector space is induced by an element of the group.

In particular, the group acts transitively on points, on collinear pairs of points, and on non-collinear pairs of points.

So the *graph* of the polar space (whose vertices are the points, two vertices joined if they are collinear) is a rank 3 graph.

In the case $r = 1$, there are no lines, so the graph of the polar space is null; Witt's lemma implies that the action of the group is 2-transitive. We will ignore this case.

6.2.5. An example. The polar space of type $\mathrm{P}\Omega^{+}(4,q)$ is the familiar *ruled quadric*, see Figure 10.

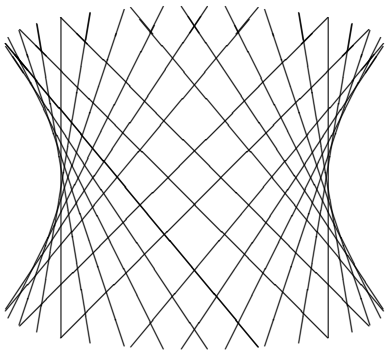

Figure 10. A ruled quadric

Combinatorially this structure is just a grid, so the classical group is non-basic. We will also ignore this case.

6.2.6. Cliques and cocliques. We must now look at cliques and cocliques in the graph Γ of a polar space.

A clique is a set of 1-dimensional subspaces on which the form vanishes and which are pairwise orthogonal; so its span is also a clique. Thus the cliques of maximal size are just the maximal subspaces, of size $(q^r-1)/(q-1)$.

Hence a coclique contains at most $q^{r+\epsilon}+1$ points, with equality if and only if it meets every maximal clique in exactly one point.

A coclique meeting this bound is called an *ovoid*.

We need one further definition: a *spread* is a family of maximal subspaces which partitions the set of points.

Theorem 6.8. (a) *A classical group is non-separating if and only if its polar space possesses an ovoid.*

(b) *A classical group is non-synchronizing if and only if its polar space possesses either*

(i) *an ovoid and a spread; or*

(ii) *a partition into ovoids.*

(c) *A classical group is not partition-separating if and only if its polar space possesses both a spread and a partition into ovoids.*

Proof. There is just one complementary pair of G-invariant graphs, for a classical group G acting on its polar space: the graph Γ defined above and its complement. We saw before the statement of the theorem that the product of the clique numbers of these graphs is equal to the number of vertices if and only if an ovoid exists. Moreover, a partition into cliques of maximal size in Γ is a spread; a partition into independent sets can only have the number of parts equal to the size of a maximal clique if there is a partition into ovoids. So the theorem follows from Corollary 4.5 and Theorem 5.4. □

6.2.7. Ovoids, spreads and partitions. You might expect at this point to be told that the question of which polar spaces contain ovoids, spreads, or partitions into ovoids has been completely solved by finite geometers.

Unfortunately, despite a lot of effort, this is not the case.

We summarize some of the results which have been obtained. A comprehensive up-to-date survey is given in [63, Chapter 7], which gives more details than we have included here. See also [45].

Ovoids

$\mathrm{PSp}(2r, q)$	Yes for $r = 2$ and q even; no in all other cases
$\mathrm{PSU}(2r, q_0)$	Yes for $r = 2$; no for some values of q_0 for $r > 2$
$\mathrm{PSU}(2r+1, q_0)$	No
$\mathrm{P}\Omega^+(2r, q)$	Yes for $n = 2, 3$; yes for $r = 4$ and q a power of 2 or 3, or q prime, or $q \equiv 5 \bmod 6$; no for $r > 4$ in some cases (for details please check the references)
$\mathrm{P}\Omega(2r+1, q)$	Yes for $r = 2$; yes for $r = 3$ and $q = 3^h$; no for $r \geq 4$
$\mathrm{P}\Omega^-(2r+2, q)$	No

Spreads

$\mathrm{PSp}(2r,q)$	Yes
$\mathrm{PSU}(2r,q_0)$	No
$\mathrm{PSU}(2r+1,q_0)$	No for $r=2$, $q_0=2$
$\mathrm{P\Omega}^+(2r,q)$	No if r is odd; yes if $r=2$, or $r=4$ with q prime or $q\equiv 3$ or 5 mod 6; yes if r and q are even
$\mathrm{P\Omega}(2r+1,q)$	No if r is even (and q odd); yes if $r=3$ with q prime or $q\equiv 3$ or 5 mod 6
$\mathrm{P\Omega}^-(2r+2,q)$	Yes if $r=2$, or if q is even

6.2.8. Some conclusions. We conclude that $\mathrm{PSp}(2r,q)$, $\mathrm{PSU}(2r+1,q_0)$, and $\mathrm{P\Omega}^-(2r+2,q)$ are separating for all $r\geq 2$, except for $\mathrm{PSp}(4,q)$ with q even. Cases where the group is not separating can also be read off from the first table. However, less is known about partitions into ovoids, so results about synchronization and partition separation are less clear. Work on this has begun recently, partly as a result of the application to synchronization.

Example 6.9. The polar space of the group $\mathrm{P\Omega}(5,q)$, for q odd, possesses ovoids but no spreads. It is proved in [17] that, if q is an odd prime, then ovoids in this polar space are all *classical*; that is, they consist of the set of points lying in a non-singular 4-dimensional space of type $\mathrm{P\Omega}^-(4,q)$ (this polar space has Witt index 1, so contains no lines). Any two such spaces meet in a 3-dimensional space, so two such ovoids meet in a *conic*. In particular, there are no partitions into ovoids.

So the group $\mathrm{P\Omega}(5,q)$, for q an odd prime, is synchronizing but not separating. These are our first examples of such groups, and show that the implication from separating to synchronizing does not reverse.

6.2.9. Spreading. We give a necessary condition for a classical group to be non-spreading, which applies to three of the six types.

Theorem 6.10. *Let G be a classical group of Witt index at least 2, acting on the points of its polar space. Suppose that there exists a non-degenerate hyperplane of the underlying vector space on which the form has Witt index smaller than that of the whole space. Then G is non-spreading.*

Proof. We take A to be a maximal subspace, and B to be the set of points lying in the assumed hyperplane. Then $|A\cap B^g| = (q^{r-1}-1)/(q-1)$ for all $g\in G$, and A and B are both sets with $|A|$ dividing $|\Omega|$. □

This theorem covers the classical groups $\mathrm{PSU}(2r,q_0)$, $\mathrm{P\Omega}^+(2r,q)$, and $\mathrm{P\Omega}(2r+1,q)$, but not $\mathrm{PSp}(2r,q)$, $\mathrm{PSU}(2r+1,q_0)$, or $\mathrm{P\Omega}^-(2r+2,q)$.

6.3. S_{2m} on (m,m) partitions. As noted earlier, it seems that testing any class of primitive groups for synchronization will produce difficult combinatorial problems. We are going to prove one more result in this section, concerning the (primitive) action of the symmetric group of even degree $2m$ on partitions of the domain into two sets of size m. This is partly because the fact that this group is non-spreading would follow from the truth of the *Hadamard conjecture*, and also because of an unexpected appearance of the *Catalan numbers* in the proof. The Catalan numbers (C_n) form one of the most ubiquitous integer sequences in all mathematics [96], but we only need two simple properties of them:

- the formula: $C_n = \dfrac{1}{n+1}\binom{2n}{n}$;
- the recurrence relation: $C_m = \sum_{i=1}^{n-1} C_i C_{n-i}$ for $i > 1$.

The Catalan numbers arise in a technical result we need. Note that the number of (m,m) partitions of a $2m$-set is $\frac{1}{2}\binom{2m}{m}$.

Lemma 6.11. *For any positive integer m,*

(a) *$2m-1$ divides $\frac{1}{2}\binom{2m}{m}$;*

(b) *if m is odd then $2(2m-1)$ divides $\frac{1}{2}\binom{2m}{m}$.*

Proof. We have

$$\frac{1}{2}\binom{2m}{m} = \binom{2m-1}{m-1} = (2m-1)\frac{(2m-2)!}{m!(m-1)!} = (2m-1)C_{m-1}.$$

If m is odd, then $m-1$ is even and the terms in the recurrence for C_{m-1} come in equal pairs. □

A *Hadamard matrix* of order n is an $n \times n$ matrix H with entries ± 1 satisfying $HH^\top = nI$. These matrices are so-called because they attain equality in Hadamard's bound for the determinant of a square matrix $A = (a_{ij})$ with $|a_{ij}| \le 1$ for all i, j.

The defining condition shows that any two rows of H are orthogonal. But it follows that $H^\top H = nI$, and so any two columns are orthogonal.

It is known that the order of a Hadamard matrix must be 1, 2 or a multiple of 4; the *Hadamard conjecture* asserts that they exist for all such orders. This is known to be true for $n < 668$ (the last value to be resolved was $n = 428$ in 2005, [73]).

Theorem 6.12. *Suppose there exists a Hadamard matrix of order $n = 4k$. Then*

(a) *S_{4k}, acting on $(2k, 2k)$ partitions, is non-spreading;*

(b) *if k is odd, then S_{2k}, acting on (k,k) partitions, is non-spreading.*

Proof. Let H be a Hadamard matrix of order $4k$.

(a) We can normalize by changing signs of columns so that the first row of H consists entirely of $+1$ entries. Then any further row has $2k$ $+1$s and $2k$ -1s, and so defines a $(2k, 2k)$ partition. Let A be the set of these partitions. Note that $|A| = 4k - 1$. Let B be the set of all $(2k, 2k)$ partitions such that the elements 1 and 2 belong to the same part. Since the any columns of H are orthogonal, $|Ag \cap B| = 2k - 1$ for any permutation g of the columns. Finally, the lemma shows that $|A|$ divides the number of partitions. So S_{4k} on $(2k, 2k)$ partitions is non-spreading.

(b) It is a well-known fact about Hadamard matrices that any three rows of a Hadamard matrix of order $4k$ agree in k positions. (This can be found in the final part of [110].) Normalize the first row as above, and consider the set of $2k$ positions where the second row has entries $+1$; then any further row has $+1$s in k of these positions and -1 in k positions. This gives us a set A of $2(2k - 1)$ partitions of a $2k$-set of type (k, k). Exactly as above, with B the set of all (k, k) partitions where 1 and 2 lie in the same part, we find that $|Ag \cap B| = 2k - 2$ for any permutation g. The second part of the lemma shows that $|A|$ divides the number of partitions if k is odd. □

Corollary 6.13. *If the Hadamard conjecture is true, then S_{2m} acting on the set of all (m, m) partitions is non-spreading for all $m > 1$.*

6.4. Factorizations of simple groups. For a final example, we turn to the simplest diagonal primitive groups, those of the form $S \times S$, where S is a simple group, acting on S by the rule

$$(g, h) : x \mapsto g^{-1}xh$$

for $(g, h) \in S \times S$, $x \in S$. We cannot prove much here: the results are mostly descriptive.

Suppose that G is a group of this form. Then questions about synchronization and separation in G reduce to questions about subsets and partitions of the simple group S. Consider the case when S is not separating, so that there exist subsets A, B of S such that $|g^{-1}Ah \cap B| = 1$ for all $g, h \in S$.

Consider first the case where A and B are subgroups of S. Then $A \cap B = 1$ and $AB = S$, so we have a *perfect factorization* of S. Conversely, suppose that we have a perfect factorization of S, and let $g = a_1b_1$ and $h = a_2b_2$, where $a_1, a_2 \in A$ and $b_1, b_2 \in B$. Then

$$\begin{aligned} |g^{-1}Ah \cap B| &= |b_1^{-1}a_1Aa_2b_2 \cap B| \\ &= |b_1^{-1}Ab_2 \cap B| \\ &= |A \cap b_1Bb_2^{-1}| \\ &= |A \cap B| \\ &= 1, \end{aligned}$$

so G is not separating.

Moreover, in this case, every right coset of A intersects every left coset of B in a single element; so the partitions of S into right cosets of A and left cosets of B demonstrate that S is not partition-separating, and so also not synchronizing.

In a perfect factorization of S, if we take the action of S on the set of right cosets of B, then A is a regular subgroup, and *vice versa*. So finding all perfect factorizations with one factor maximal is equivalent to finding all regular subgroups of primitive groups (since a permutation group is primitive if and only if the point stabiliser is a maximal subgroup). This problem has been solved by Liebeck, Praeger and Saxl (see [79]).

In the case where one of A and B is a subgroup (say A) and the other is not, the condition that A and B witness the non-separating property of G is equivalent to saying that B is a *loop transversal* for A in S, so that in the action of S on the right cosets of A, the set B is sharply transitive: see [66, 75], for example.

We do not know of any examples where neither A nor B is a subgroup, though no doubt they exist.

7. Representation theory

The concept of "spreading" defined earlier turns out to be expressible in terms of representation theory. In this section we outline the permutation representation of a permutation group, and show how its properties over different fields are related to some of the concepts we are considering.

7.1. 2-closure. We have seen that synchronization and related properties are closed upwards (i.e. preserved on passing to overgroups). They also have a limited form of downward closure, as we will now see.

Let G be a permutation group on Ω.

(a) The 2-*closure* of G is the set of all permutations of Ω which preserve the G-orbits on Ω^2 (the set of ordered pairs of elements of Ω). The group G is 2-*closed* if it is equal to its 2-closure.

(b) The *strong* 2-*closure* of G is the set of all permutations of Ω which preserve the G-orbits on the set of 2-element subsets of Ω. The group G is *strongly* 2-*closed* if it is equal to its strong 2-closure.

Note that

(a) the 2-closure of G is contained (possibly strictly) in its strong 2-closure;

(b) the 2-closure of G is the symmetric group if and only if G is 2-transitive;

(c) the strong 2-closure of G is the symmetric group if and only if G is 2-homogeneous.

Theorem 7.1. *Let P denote one of the conditions "primitive", "synchronizing", "separating", "2-homogeneous". Then the following are equivalent:*

(a) *G satisfies P;*

(b) *the 2-closure of G satisfies P;*

(c) *the strong 2-closure of G satisfies P.*

Proof. In view of our earlier remarks, (a) implies (b) implies (c); so it suffices to show that (c) implies (a). But each property can be defined in terms of G-invariant graphs, and G and its strong 2-closure clearly preserve the same graphs. □

7.2. Representation theory. We now turn to an algebraic approach to these and related closure properties. Let $\mathbb{F}$ be a field. We only consider the case $\mathbb{F} = \mathbb{C}$, $\mathbb{R}$ or $\mathbb{Q}$. Certainly there is an interesting theory waiting to be worked out in the case where $\mathbb{F}$ is, say, a finite field, a p-adic field, or even a ring!

Let G be a permutation group on Ω. The *permutation module* is the $\mathbb{F}G$-module $\mathbb{F}\Omega$ which has the elements of Ω as a basis, where G acts by permuting the basis vectors.

Now the $\mathbb{F}$-*closure* of G consists of all permutations which preserve all $\mathbb{F}G$-submodules of $\mathbb{F}\Omega$; and G is $\mathbb{F}$-*closed* if it is equal to its $\mathbb{F}$-closure.

Consider the case where G is the symmetric group $\mathrm{Sym}(\Omega)$. The permutation module has just two non-trivial submodules:

(a) the 1-dimensional module $\underline{\Omega}$ spanned by the sum of the elements of Ω;

(b) the $n-1$-dimensional *augmentation submodule* consisting of the vectors with coordinate sum zero.

For, if W is a submodule containing a vector x with $x_v \neq x_w$, and g is the transposition (v, w), then W contains $x - xg = \lambda(v - w)$. By 2-transitivity, W contains all differences between basis vectors; but these span the augmentation module.

Theorem 7.2. *The $\mathbb{C}$-closure of a permutation group G is equal to its 2-closure.*

The proof requires a little character theory; a brief sketch follows.

7.3. Character theory. Any representation of a group by matrices over the complex numbers is determined up to isomorphism by its *character*, the function ϕ which maps each group element to the trace of the matrix representing it. A character is a *class function* (constant on conjugacy classes).

Any representation can be decomposed uniquely (up to isomorphism) into *irreducible* representations. An *irreducible character* is the character of an irreducible representation.

The irreducible characters form an orthonormal basis for the space of complex class functions, under the inner product

$$\langle \phi, \psi \rangle = \frac{1}{|G|} \sum_{g \in G} \phi(g)\overline{\psi(g)}.$$

The *trivial character* 1_G is the function mapping every group element to 1.

7.4. The permutation character. Let G be a permutation group on Ω, where $|\Omega| = n$. Then we have an action of G on $\mathbb{C}\Omega$ by permutation matrices. Its character is the *permutation character* π, where $\pi(g)$ is the number of fixed points of g.

The *Orbit-Counting Lemma* states that

$$\frac{1}{|G|} \sum_{g \in G} \pi(g) = \#\text{ orbits of } G.$$

The sum on the left is just $\langle 1_G, \pi \rangle$; so the multiplicity of the trivial character in π is equal to the number of orbits of G.

Applying the preceding result to the action of G on $\Omega \times \Omega$ (whose permutation character is π^2), we see that

$$\langle \pi, \pi \rangle = \langle \pi^2, 1_G \rangle = \#\text{ orbits of } G \text{ on } \Omega^2.$$

This number is called the *rank* of G.

The rank is equal to the sum of squares of the multiplicities of the irreducible characters in π, since if $\pi = \sum a_i \phi_i$, with ϕ_i irreducible, then orthonormality gives

$$\langle \pi, \pi \rangle = \sum a_i^2.$$

In particular, G is 2-transitive if and only if $\pi = 1_G + \phi$ for some irreducible character ϕ. (The character ϕ is afforded by the action of G on the augmentation submodule of the permutation module: so G is 2-transitive if and only if the augmentation submodule is irreducible.)

We recall that orbits of G on $\Omega \times \Omega$ are called *orbitals* of G. Given an orbital O, there is a *paired* orbital

$$O^* = \{(y, x) : (x, y) \in O\}.$$

An orbital O is *self-paired* if $O = O^*$. We see from the above that if G has permutation character $\pi = \sum a_\phi \phi$, where ϕ are irreducible characters of G, then the number of orbitals is $\sum a_\phi^2$.

We will discuss further the combinatorial structure of the orbitals in Section 10.1.

The decomposition of the permutation character also tells us about the number of self-paired orbitals. This involves the *Frobenius–Schur index* ϵ_ϕ of an irreducible character ϕ, defined as follows:

$$\epsilon_\phi = \begin{cases} 1 & \text{if } \phi \text{ is the character of a real representation of } G, \\ -1 & \text{if } \phi \text{ is real-valued but not the character of a real representation,} \\ 0 & \text{if } \phi \text{ is not real-valued.} \end{cases}$$

A character ϕ is called *real*, *quaternionic* or *complex* according as $\epsilon_\phi = +1, -1$ or 0. (The term refers to the centralizer algebra of the corresponding real representation affording the character ϕ, 2ϕ, or $\phi + \overline{\phi}$ respectively.)

Theorem 7.3. *Let the permutation character of G be*

$$\pi = \sum_\phi a_\phi \phi,$$

where ϕ are irreducible characters of G. Then the number of self-paired orbitals of G is $\sum \epsilon_\phi a_\phi$.

An elementary account of this theorem appears in [33].

7.5. 2-closure = $\mathbb{C}$-closure. We are now ready to prove Theorem 7.2.

Let $\bar{G}$ be the 2-closure of G. Then $\bar{G}$ has the same sum of squares of multiplicities of irreducibles as G, since this sum of squares is equal to the number of orbits on ordered pairs. This implies that the decomposition of the permutation character is the same for $\bar{G}$ as for G. (If a character becomes reducible on restriction, then either the number of irreducibles or the multiplicity of at least one will increase.) Hence $\bar{G}$ is contained in the $\mathbb{C}$-closure of G.

Conversely, let $\hat{G}$ be the $\mathbb{C}$-closure of G. Then $\hat{G}$ preserves the isotypic components of the permutation module (one of these consists of the sum of all copies of a particular isomorphism type of irreducible module). The lattice of submodules of the sum of r isomorphic irreducible modules is isomorphic to the $(r-1)$-dimensional

complex projective space; all these submodules are preserved by $\hat{G}$. So the isomorphic G-modules remain isomorphic as $\hat{G}$-modules. Thus the multiplicities are the same for $\hat{G}$ as for G, and so the ranks of these groups are equal. Since $G \leq \hat{G}$, it follows that $\hat{G}$ preserves the G-orbits on Ω^2, and so is contained in the 2-closure $\bar{G}$.

Hence $\hat{G} = \bar{G}$.

Conjecture 7.4. The $\mathbb{R}$-closure of a permutation group coincides with its strong 2-closure.

This is not known in general, but it is true for groups whose permutation character is multiplicity-free.

7.6. $\mathbb{F}$I groups. We say that the permutation group G on Ω is $\mathbb{F}$I if its $\mathbb{F}$-closure is the symmetric group; that is, if the only G-submodules of $\mathbb{F}\Omega$ are $\underline{\Omega}$ and the augmentation module. Note that if $\mathbb{F} \subseteq \mathbb{K}$, then $\mathbb{K}$I implies $\mathbb{F}$I because extension of scalars commute with direct sums.

Theorem 7.5. *Let G be a permutation group on Ω.*

(a) *G is $\mathbb{C}$I if and only if it is 2-transitive.*

(b) *G is $\mathbb{R}$I if and only if it is 2-homogeneous.*

Proof. (a) G is $\mathbb{C}$I if its permutation character has the form $1_G + \phi$, where ϕ is irreducible over $\mathbb{C}$. As noted above, this is equivalent to the assertion that G is 2-transitive.

(b) G is $\mathbb{R}$I if its permutation character has the form $1_G + \theta$, where θ is irreducible over $\mathbb{R}$. Now there are three possibilities for the decomposition of θ over $\mathbb{C}$:

- θ is irreducible over $\mathbb{C}$: then G is 2-transitive by the preceding argument.
- $\theta = 2\phi$, where ϕ is irreducible over $\mathbb{C}$. Then $\epsilon_\phi = -1$, and so the number of self-paired orbitals of G is $1 - 2 = -1$, which is impossible.
- $\theta = \phi + \overline{\phi}$ for some non-real-valued character ϕ. Then $\epsilon_\phi = 0$, and so the number of self-paired orbitals is 1, this one being the diagonal orbital. Thus the two non-diagonal orbitals are paired with each other, and G is 2-homogeneous.

The argument clearly reverses.

The non-existence in the second case also follows from an old result of Jordan (see Serre [95]), according to which a finite transitive permutation group of degree greater than 1 contains a fixed-point-free element. Now, if $\pi = 1_G + 2\phi$ and $\pi(g) = 0$, then $\phi(g) = -\frac{1}{2}$, contradicting the fact that character values must be algebraic integers. □

This naturally suggests looking at $\mathbb{Q}$I groups, to which we now turn.

Theorem 7.6. *Let G be a transitive permutation group on Ω, and $\mathbb{F}$ a field of characteristic zero. Then G is primitive (resp. synchronizing, separating, spreading, or $\mathbb{Q}$I) if and only if its $\mathbb{F}$-closure is.*

Three of these results are immediate from the next lemma.

Lemma 7.7. *Let G be a transitive permutation group on Ω, and $\mathbb{F}$ a field of characteristic zero. Let A and B be multisets such that $|A*Bg| = \lambda$ for all $g \in G$. Then $|A*Bg| = \lambda$ for all g in the $\mathbb{F}$-closure $\hat{G}$ of G.*

Proof. Let v_1 and v_2 be the characteristic functions of A and B respectively. Set $w_i = v_i - (v_i \cdot j)j/n$ for $i = 1, 2$, where j is the all-1 vector. Now, using the facts that $\lambda = v_1 \cdot v_2 g = \frac{(v_1 \cdot j)(v_2 \cdot j)}{n}$ by Theorem 5.12, and $j \cdot j = n$, we find that j, w_1 and $w_2 g$ are pairwise orthogonal for any $g \in G$. So the G-submodules generated by j, w_1 and w_2 are pairwise orthogonal. These modules are invariant under $\hat{G}$; reversing the calculations gives the result. □

This immediately proves Theorem 7.6 for separating, spreading and $\mathbb{Q}$I groups.

Suppose that G is imprimitive, and let P be a G-invariant partition. Then the characteristic functions of the parts of P form an orthogonal basis for a submodule of $\mathbb{F}\Omega$, which is preserved by $\hat{G}$. The partition can be recovered from the submodule, since it is the coarsest partition on the parts of which the elements of the submodule are constant. So $\hat{G}$ is imprimitive.

Finally, suppose that G is not synchronizing, and let the partition P and section S witness this (that is, each G-translate of S is a section of P, see Theorem 3.8). Then $|A \cap Sg| = 1$ for any part A of P and $g \in G$ and so, by the lemma, $|A \cap Sg| = 1$ for all $g \in \hat{G}$. Thus every $\hat{G}$-translate of S is a section for P and so $\hat{G}$ is non-synchronizing. This completes the proof.

If G is $\mathbb{Q}$I, then its $\mathbb{Q}$-closure is the symmetric group, which is spreading; so G is spreading. This was first proved in [14].

So our hierarchy finally looks like this:

$$\begin{aligned}
&\text{2-transitive} \Rightarrow \text{2-homogeneous} \Rightarrow \mathbb{Q}\text{I} \Rightarrow \\
&\quad \Rightarrow \text{spreading} \Rightarrow \text{separating} \Rightarrow \text{synchronizing} \Rightarrow \\
&\Rightarrow \text{basic \& almost synchronizing} \quad \Rightarrow \quad \left.\begin{array}{c}\text{basic}\\ \text{or}\\ \text{almost synchronizing}\end{array}\right\} \Rightarrow \\
&\Rightarrow \text{primitive} \Rightarrow \text{transitive.}
\end{aligned}$$

We will see that there are groups which are $\mathbb{Q}$I but not 2-homogeneous; indeed, these groups have recently been classified. But no examples are currently known of groups which are spreading but not $\mathbb{Q}$I.

Theorem 6.10 suggested to us that the classical groups $\mathrm{PSp}(2r, q)$, $\mathrm{PSU}(2r+1, q_0)$, and $\mathrm{P}\Omega^-(2r+2, q)$ may be good candidates for groups which are spreading but not $\mathbb{Q}$I. However, Pablo Spiga (private communication to the authors) was able to show that $\mathrm{PSp}(4, p)$ is non-spreading for $p = 3, 5, 7$ by computational methods. The issue is unresolved in general.

We have seen examples of basic, but not almost synchronizing; and of almost synchronizing, but non-basic; and, of course, synchronizing groups are almost synchronizing and basic.

7.7. Affine groups. Recall that an *affine group* is a permutation group G on the d-dimensional vector space over $\mathbb{F}_p$ (where p is prime) generated by the translation group T and an irreducible linear group H. Thus G is the semidirect product of T by H; and H is the stabilizer of the zero vector.

Theorem 7.8. *Let G be an affine permutation group on the d-dimensional vector space over $\mathbb{F}_p$, with $H = G_0$ as above. Then the following are equivalent:*

(a) *G is spreading;*

(b) *G is $\mathbb{Q}$I;*

(c) *H is transitive on the set of 1-dimensional subspaces of V;*

(d) *the group generated by G and the scalars in $\mathbb{F}_p$ is 2-transitive.*

The affine groups described in the Theorem can be classified, using the classification of affine 2-transitive groups [60, 77].

Proof. It is clear that (c) and (d) are equivalent. Let us suppose that they do not hold. Then H is not transitive on 1-dimensional spaces of V, and hence not transitive on $(d-1)$-dimensional subspaces either (by Brauer's lemma). Choose hyperplanes A and B in different orbits of H. Then no image of B under G is parallel to A, so $|A \cap Bg| = p^{d-2}$ for all $g \in G$. Thus G is not spreading. So (a) implies (c) and (d). It is clear that (b) implies (a); so it remains to prove that (c) and (d) imply that G is $\mathbb{Q}$I.

The scalars in $\mathbb{F}_p$ act on V, and hence on the characters of V; and their action is precisely that of the Galois group of the field of pth roots of unity. Now assuming that (d) holds, this group permutes the non-principal irreducibles in the permutation character transitively, and so G is $\mathbb{Q}$I, as required. □

The equivalence of (b)–(d) is due to Dixon [46].

A transitive group G of prime degree p contains a p-cycle a. Clearly (c) of the theorem is satisfied by the affine group $C = \langle a \rangle$ and so C, and hence its overgroup G, is spreading. Also, each element of C can be expressed by a word of length at most $p-1$ in $\{a\}$. Therefore, if f is any singular mapping, then Theorem 5.16 yields a reset word over $\{a, f\}$ of length at most $1+p(p-2) = (p-1)^2$. This result, due to Pin [86], was the first positive result concerning the Černý conjecture.

7.8. 3/2-transitive groups. A permutation group G on Ω is said to be 3/2-*transitive* if it is transitive, and the stabilizer of a point v has all its orbits except $\{v\}$ of the same size. (If there is just one such orbit then G is 2-transitive.)

Example 7.9. Let q be a power of 2. The group $\mathrm{PSL}(2, q)$ has dihedral subgroups of order $2(q+1)$; it acts transitively on the set of cosets of such a subgroup, and the stabilizer has $q/2-1$ orbits each of size $q+1$ on the remaining points.

Example 7.10. There is a "sporadic" example: the symmetric group S_7 acting on 2-subsets of $\{1, \ldots, 7\}$. This works because $2 \cdot 5 = \binom{5}{2}$.

Using the Classification of Finite Simple Groups, John Bamberg, Michael Giudici, Martin Liebeck, Cheryl Praeger and Jan Saxl [18] have determined the almost simple 3/2-transitive groups. Apart from the affine groups of Theorem 7.8, the only primitive 3/2-transitive groups are those of Examples 7.9 and 7.10.

Although the class of 3/2-transitive groups is not closed upwards, this classification gives us the $\mathbb{Q}$I-groups:

Theorem 7.11. *Any* $\mathbb{Q}$I *group is* 3/2*-transitive.*

The reason is that the permutation character is the sum of the trivial character and a family of algebraically conjugate characters; an old result of Frame [52, 53] (see [113, §30]) now applies. This was first observed by Dixon [46].

Now the group S_7 acting on 2-sets is not $\mathbb{Q}$I. Careful analysis of the character values of $\mathrm{PSL}(2,q)$ show that the 3/2-transitive action of this group described earlier is $\mathbb{Q}$I if and only if $q-1$ is a Mersenne prime. So there are probably infinitely many examples of this form (the *Lenstra–Pomerance–Wagstaff conjecture* [109]), though nobody knows for sure.

Any other $\mathbb{Q}$I group is affine, and as we have seen in Theorems 7.8 and 7.11, these groups are classified up to the existence of Mersenne primes.

We conclude this section by showing that, for some special degrees, our hierarchy of properties collapses.

Theorem 7.12. *Let* G *be a permutation group of degree* n*.*

(a) *If* n *is prime, then* G *is* $\mathbb{Q}$I *if and only if it is transitive.*

(b) *If* n *is the square of a prime, then* G *is* $\mathbb{Q}$I *if and only if it is synchronizing.*

Proof. (a) According to a theorem of Burnside [31], a transitive group of prime degree is either contained in the 1-dimensional affine group (so that its $\mathbb{Q}$-closure is the full affine group, which is 2-transitive), or 2-transitive.

(b) A theorem of Wielandt [114] asserts that a primitive group G of degree p^2, where p is prime, satisfies one of the following: G is affine; G is contained in S_p wr S_2; or G is 2-transitive. The third case requires no further comment, and in the second case G is non-basic and hence non-synchronizing. So consider the case when G is a subgroup of $\mathrm{AGL}(2,p)$. Let H be the stabiliser of the origin, so that $H \leq \mathrm{GL}(2,p)$. If H acts transitively on the set of 1-dimensional subspaces of the 2-dimensional vector space, then G is $\mathbb{Q}$I, by Theorem 7.8. Otherwise, take two subspaces U and W in different H-orbits; it is easy to see that all G-images of U are sections for the partition into cosets of W. (This argument shows that in fact we could replace "synchronizing" by "partition-separating" in the statement of the theorem.) □

It is natural to wonder whether the synchronization hierarchy simplifies for other special degrees for which a classification of the primitive groups exist. These include squarefree integers [76] and prime powers [32]. However, examining the lists in these papers shows that they contain many examples of primitive groups for which synchronization involves difficult combinantorial or geometric problems, such as S_n on k-sets in the first case, and affine groups in the second. There is

one case in which we can reach a conclusion: If $n = 2p$, where p is prime, then a primitive group of degree n is synchronizing (see [84, Corollary 2.5]).

To put these results in context, it may be worth mentioning the following result [39]:

Theorem 7.13. *For almost all natural numbers n (all but a set of density zero), the only primitive groups of degree n are the symmetric and alternating groups.*

7.9. $\mathbb{Q}$I versus spreading. We don't know any examples of groups which are spreading but not $\mathbb{Q}$I. Moreover, there are very few $\mathbb{Q}$I groups, and there are plenty of places to look for spreading groups.

We have

- G is not $\mathbb{Q}$I if and only if there are non-trivial multisets A and B satisfying $(1)_\lambda$.

For, if G fails to be $\mathbb{Q}$I, take two orthogonal submodules of the augmentation module, and choose non-zero vectors w_1, w_2 in these submodules. Multiplying up by a suitable integer, we can assume that the coordinates of w_1 and w_2 are integers; adding a suitable multiple of j, we can assume that the coordinates are non-negative, and so the resulting vectors are] multisets satisfying $(1)_\lambda$ for some λ. The argument reverses (as in the proof of Lemma 7.7).

On the other hand, by definition (see Subsection 5.5), we have

- G is not spreading if and only if there are non-trivial multisets A and B satisfying $(1)_\lambda$, (3) and (4).

Condition (3) says that B is a set. In combinatorial problems of this kind, there is usually a big difference between asking for a multiset with a certain property and asking for a set. This reason and others suggest that such groups will exist; but none have yet been found!

Further applications of representation theory to synchronization can be found in [4, 101].

8. Detecting properties with functions

In this section we are going to expand on the detection of the primitive, synchronizing, spreading or separating properties using functions. The first, motivating result is Rystsov's Theorem (Theorem 2.8), which we reformulate here:

Theorem 8.1. *Let G be a transitive permutation group on Ω, with $|\Omega| = n$. The following are equivalent:*

(a) *G is primitive;*

(b) *for any function $f\colon \Omega \to \Omega$ whose image has cardinality $n-1$, the semigroup generated by G and f contains a constant function;*

(c) *for any idempotent function $f\colon \Omega \to \Omega$ whose image has cardinality $n-1$, the semigroup generated by G and f contains a constant function.*

Rystsov does not, in fact, explicitly state the above theorem. But in [91] he proved that if $a \neq b$ are elements of Ω, then the orbital digraph corresponding to (a, b) is connected if and only if G synchronizes the rank $n-1$ idempotent e defined by

$$xe = \begin{cases} x & \text{if } x \neq a; \\ b & \text{if x=a.} \end{cases}$$

Theorem 8.1 is an immediate consequence of this result, Higman's characterization of primitivity in terms of the connectivity of orbital digraphs and the easy observation that if G is a transitive group and f is any rank $n-1$ mapping, then $\langle G \cup f \rangle$ contains a rank $n-1$ idempotent.

Theorem 8.2. *Let G be a permutation group on Ω, with $|\Omega| = n$. The following are equivalent:*

(a) *G is 2-homogeneous;*

(b) *for any function $f\colon \Omega \to \Omega$ whose image has size $n-1$, the semigroup generated by G and f contains all transformations which are not permutations;*

(c) *for any idempotent function $f\colon \Omega \to \Omega$ whose image has size $n-1$, the semigroup generated by G and f contains all transformations which are not permutations.*

Proof. It is obvious that (b) implies (c).

To prove that (a) implies (b), let G be 2-homogeneous and let $f\colon \Omega \to \Omega$ be a rank $n-1$ map, whose unique non-singleton kernel class is $\{a, b\}$, and $\{a_0\} = \Omega \setminus \Omega f$. We claim that $\langle G, f \rangle$ contains all idempotents of rank $n-1$. In fact, if e is one such idempotent, with non-singleton kernel class $\{c, d\}$ and $\{c_0\} = \Omega \setminus \Omega e$, there exist $h, g \in G$ such that $\{c, d\}g = \{a, b\}$ and $a_0 h = c_0$. Therefore, the unique non-singleton kernel class of gfh is $\{c, d\}$ and $\Omega gfh = \Omega e$; so gfh has the same image and kernel as e. Since e is idempotent, it follows that its image is a section of its kernel, and hence the same happens with gfh; thus $\mathrm{rank}(gfh)^k = \mathrm{rank}(e) = n-1$, for all natural numbers k; now we can partition the natural number as follows: for

all natural i and j, we say that $i \sim j$ if $(gfh)^i = (gfh)^j$. By Schur's Theorem [94], there exists a part in this partition containing a, b and $a+b$, that is,

$$(gfh)^a = (gfh)^{a+b} = (gfh)^a(gfh)^b = (gfh)^a(gfh)^a = ((gfh)^a)^2,$$

and hence there exists a natural number a such that $(gfh)^a$ is idempotent, having the same kernel and image as e. But this forces $(gfh)^a = e$ (because there is a unique idempotent with given image and kernel), and the claim is proved. It is well known, see [64], that the rank $n-1$ idempotent maps generate all non-invertible maps and this concludes the proof of the implication.

To prove that (c) implies (a) suppose e is an idempotent with non-singleton kernel class $\{a, b\}$ such that $\langle f, G\rangle$ contains all non-invertible maps. Let $f' \colon \Omega \to \Omega$ be a rank $n-1$ map with non-singleton kernel class $\{c, d\}$. Since, by hypothesis, $f' \in \langle e, G\rangle$, it follows that $f' = g_1 e g_2 \dots e g_k$, and hence $(c, d) \in \ker(g_1 e g_2 \dots e g_k)$. As $\operatorname{rank}(g_1 e g_2 \dots e g_k) = \operatorname{rank}(e)$ it follows that $\ker(g_1 e g_2 \dots e g_k) = \ker(g_1 e)$; thus there exists $g_1 \in G$ such that $\{c, d\}g_1 = \{a, b\}$. As $\{c, d\}$ was an arbitrary 2-set, it follows that G has only one orbit on 2-sets. The implication follows. □

The above theorem was first proved by McAlister [81].

Denote by $\operatorname{Unif}(\Omega)$ the set of functions $f \colon \Omega \to \Omega$ whose kernel is a uniform partition with at least two parts. The next result provides some characterizations of synchronizing groups.

Theorem 8.3. *Let G be a transitive permutation group on Ω, with $|\Omega| = n$. The following are equivalent:*

(a) *there is no non-trivial partition P and set A such that Ag is a section for P, for all $g \in G$;*

(b) *for any function $f \colon \Omega \to \Omega$ which is not a permutation, the semigroup generated by G and f contains a constant function;*

(c) *for any idempotent function $f \colon \Omega \to \Omega$ which is not a permutation, the semigroup generated by G and f contains a constant function;*

(d) *for all $t \in \operatorname{Unif}(\Omega)$, there exists a part A of $\operatorname{Ker}(t)$ and $g \in G$ such that $|\Omega t g \cap A| > 1$.*

Proof. The equivalence (a) and (b) is the content of Theorem 3.8; the equivalence of (b) and (c) is immediate since every mapping has an idempotent positive power.

The implication (d) implies (b) is essentially the content of Corollary 3.10: if G is not synchronizing, then a minimal rank mapping t not synchronized by G belongs to $\operatorname{Unif}(\Omega)$. By (d), there exist $g \in G$ and a part A of $\operatorname{Ker}(t)$ such that $|\Omega t g \cap A| > 1$. Therefore $\operatorname{rank}(tgt) < \operatorname{rank}(t)$, a contradiction to the choice of t.

Conversely, suppose that $t \in \operatorname{Unif}(\Omega)$ is such that $\langle G, t\rangle$ contains a constant mapping. Then there exists $g \in G$ such that $\operatorname{rank}(tgt) < \operatorname{rank}(t)$, which is equivalent to saying that, for some part A of $\operatorname{Ker}(t)$, we have $|\Omega t g \cap A| > 1$. The result follows. □

The following result provides a characterization of separation that parallels the equivalence of (a) and (d) in the previous result.

Theorem 8.4. *Let G be a transitive permutation group on Ω, with $|\Omega| = n$. The following are equivalent:*

(a) *G is separating;*

(b) *for all $t \in \mathrm{Unif}(\Omega)$ and all parts A of $\mathrm{Ker}(t)$, there exists $g \in G$ such that $|A \cap \Omega tg| > 1$.*

Proof. Suppose that G is separating, and let $t \in \mathrm{Unif}(\Omega)$ be singular. Put $B = \Omega t$ and let A be an arbitrary $\mathrm{Ker}(t)$-class. Then $|A| \cdot |B| = |\Omega|$ (because t is uniform) and there exists $g \in G$ such that $|A \cap Bg| > 1$, as G is separating and the average value of $|A \cap Bg|$ is 1 by Theorem 5.12.

Conversely, suppose that G is not separating. Then we have two non-trivial subsets A, B of Ω such that $|A| \cdot |B| = |\Omega|$ and $|A \cap Bg| = 1$ for all $g \in G$. Let $t \in \mathrm{Unif}(\Omega)$ be any mapping such that A is a part of $\mathrm{Ker}(t)$ and $\Omega t = B$. Then, by (b), there exists $g \in G$ such that $|A \cap Bg| > 1$, a contradiction. □

The above theorem, in light of Theorem 8.3, provides another proof that separating groups are synchronizing.

Theorem 5.13 gave a characterization of spreading groups. We close this section with another.

Theorem 8.5. *Let $G \leq S_n$ be a transitive group acting on Ω. The following are equivalent:*

(a) *for all proper subsets A of Ω and singular mappings $t \in T(\Omega)$, there exists $g \in G$ such that $|Agt^{-1}| > |A|$;*

(b) *for all proper subsets A of Ω and idempotent singular mappings $e \in T(\Omega)$, there exists $g \in G$ such that $|Age^{-1}| > |A|$.*

Proof. That (a) implies (b) is obvious. Conversely, let t be a singular mapping on Ω. Then there exists a singular idempotent mapping e and $h \in S_n$ such that $h^{-1}e = t$. By (b), there exists $g \in G$ such that $|Age^{-1}| > |A|$. Therefore, $|Agt^{-1}| = |Age^{-1}h| = |Age^{-1}| > |A|$, as required. □

Observe that synchronizing groups can be defined in terms of functions or in terms of section-regular partitions. The previous result, allowing a definition of spreading groups in terms of idempotents, leads to a parallel definition in terms of partitions and sections. A group G is spreading if and only if, for every $A \subsetneq \Omega$, and every partition P of Ω with section B, there exists $g \in G$ such that $|Ag * B'| > |A|$, where B' is the multiset with support B that gives $x \in B$ multiplicity the size of its part in P.

9. Applications to the Černý Conjecture

So far, with the exception of Theorem 5.16, we have not provided any bounds on lengths of reset words. In this section we prove a new result, generalizing previous results of Rystsov [89] and the third author [97], giving bounds for reset words in the case of a transitive permutation group and a collection of singular transformations that it synchronizes.

9.1. Transitive permutation groups. Let us set up some notation. If G is a finite group and A is a generating set, then we write $d_A(G)$ for the smallest integer $d \geq 0$ such that $G = (A \cup \{1\})^d$. One can think of $d_A(G)$ as the directed diameter of the Cayley digraph of G with respect to A. All our bounds are based on this parameter. Trivially, $d_A(G) \leq |G| - 1$ since a shortest directed path in the Cayley digraph of G with respect to A from 1 to any vertex has length at most $|G| - 1$.

If $M \subseteq T(\Omega)$ is a transformation monoid, then the $\mathbb{Q}$-vector space $\mathbb{Q}^\Omega$ of mappings $f\colon \Omega \to \mathbb{Q}$ is a left $\mathbb{Q}M$-module via the action defined by $(mf)(x) = f(xm)$ for $m \in M$ and $x \in \Omega$. Moreover, the subspace V_1 of constant mappings is a $\mathbb{Q}M$-submodule isomorphic to the trivial $\mathbb{Q}M$-module. Notice that the character θ of $\mathbb{Q}^\Omega$ is given by

$$\theta(m) = |\{x \in \Omega : xm = x\}|$$

and so, for a group, this is just the character of the permutation module $\mathbb{Q}\Omega$. However, for monoids there is a significant difference between the transformation module $\mathbb{Q}\Omega$ and its vector space dual $\mathbb{Q}^\Omega$; for instance, if M is synchronizing and transitive, then $\mathbb{Q}\Omega$ is a projective indecomposable right $\mathbb{Q}M$-module with the trivial module as its simple top, whereas $\mathbb{Q}^\Omega$ is an injective indecomposable left $\mathbb{Q}M$-module with the trivial module as its simple socle (see [101] or [98] for details).

If M is a monoid, V is a left $\mathbb{Q}M$-module, $C \subseteq M$ and $W \subseteq V$ is a subspace, then we denote by CW the linear span of all vectors of the form cw with $c \in C$ and $w \in W$.

Lemma 9.1. *Let G be a group generated by a set A and let V be a finite dimensional $\mathbb{Q}G$-module. Suppose that $W \subseteq V$ is a subspace. Then the equality $(A \cup \{1\})^d W = GW$ holds where $d = \dim V - \dim W$.*

Proof. Put $W_i = (A \cup \{1\})^i W$. Then we have an increasing chain

$$W_0 \subseteq W_1 \subseteq \cdots$$

of subspaces of V whose union is GW. It follows by dimension considerations that $W_i = W_{i+1}$ for some $0 \leq i \leq d$. But this means $W_i = GW$ and hence $W_d = GW$, as required. □

Now we can establish our desired synchronization bound.

Theorem 9.2. *Let G be a transitive permutation group on a set Ω of cardinality $n \geq 2$ and let A be a generating set for G. Suppose that $B \subseteq T(\Omega)$ is such that $\langle G \cup B\rangle$ is synchronizing. Then there is a reset word over $A \cup B$ of length at most*

$$1 + (n - m + d_A(G))(n - 2)$$

where m is the maximum dimension of an irreducible $\mathbb{Q}G$-submodule of $\mathbb{Q}^\Omega$.

In particular, in the case that $m \geq d_A(G)$, there is a reset word over $A \cup B$ of length at most $(n-1)^2$.

Proof. We claim that if $S \subsetneq \Omega$ is a proper subset with at least two elements, then there exists a word v over $A \cup B$ of length at most $n - m + d_A(G)$ such that $|Sv^{-1}| > |S|$. Let us see why this claim implies the proposition. Since $A \cup B$ is synchronizing, it contains a singular mapping $b \in B$. Let $x \in \Omega$ with $|xb^{-1}| \geq 2$. The claim then finds us a sequence $v_1, \ldots, v_k$ of words over $A \cup B$, each of length at most $n - m + d_A(G)$, such that $|xb^{-1}v_1^{-1} \cdots v_k^{-1}| = n$ with $1 \leq k \leq n - 2$. Thus $v_k v_{k-1} \cdots v_1 b$ is a reset word (with image x) of length at most $1 + (n - m + d_A(G))(n - 2)$.

To prove the claim, set $M = \langle G \cup B\rangle$ and let $V = \mathbb{Q}^\Omega$ with the left $\mathbb{Q}M$-module structure described above. There is a direct sum decomposition of V as a $\mathbb{Q}G$-module $V = V_0 \oplus V_1$ where

$$V_0 = \left\{ f \in V : \sum_{x \in \Omega} f(x) = 0 \right\}$$

is a hyperplane and V_1 is the line consisting of constant mappings. It is well known that the transitivity of G implies that V_1 is the isotypic component of the trivial $\mathbb{Q}G$-module and hence the operator $P = \sum_{g \in G} g$ annihilates V_0; indeed, P is a scalar multiple of the primitive idempotent corresponding to the trivial representation. Note that V_1 is a $\mathbb{Q}M$-submodule, but V_0 is not. We remark that m is the dimension of an irreducible $\mathbb{Q}G$-submodule of V_0.

Denote by χ_A the characteristic function of a subset $A \subseteq \Omega$ and consider the vector

$$\gamma_S = \chi_S - \frac{|S|}{n}\chi_\Omega \in V_0.$$

Let $W_0 = G\gamma_S$ be the $\mathbb{Q}G$-submodule generated by γ_S. Then $W_0 \subseteq V_0$. On the other hand, if w is a reset word over $A \cup B$ with $\Omega w \subseteq S$ (such exists by transitivity of G), then

$$w\gamma_S = \chi_{Sw^{-1}} - \frac{|S|}{n}\chi_\Omega = \left(1 - \frac{|S|}{n}\right)\chi_\Omega \notin V_0.$$

Thus $\mathbb{Q}M \cdot W_0 \nsubseteq V_0$. It follows that if W is any $\mathbb{Q}G$-submodule of V_0 containing W_0, then there exists $b \in B$ with $\{1, b\}W \supsetneq W$. Therefore, we can choose $b_1, \ldots, b_j \in B$ such that if

$$W_i = G\{1, b_i\}G\{1, b_{i-1}\}G \cdots G\{1, b_1\}W_0,$$

then we have

$$W_0 \subsetneq \cdots \subsetneq W_{j-1} \subseteq V_0$$

and $W_j \not\subseteq V_0$. For convenience, we put $W_{-1} = 0$.

Note that $\{1, b_{i+1}\}W_i \supsetneq W_i$ for $0 \leq i \leq j-1$ and so, by repeated application of Lemma 9.1,

$$W_{j-1} = (A \cup \{1\})^{d_{j-1}}\{1, b_{j-1}\} \cdots (A \cup \{1\})^{d_1}\{1, b_1\}(A \cup \{1\})^{d_0}\gamma_S$$

where $d_i = \dim W_i - \dim W_{i-1} - 1$ for $0 \leq i \leq j-1$. Thus we can find words $w_0, \ldots, w_j$ over A such that $b_j w_{j-1} b'_{j-1} \cdots w_1 b'_1 w_0 \gamma_S \notin V_0$ where $b'_i \in \{1, b_i\}$ and

$$|w_i| \leq \dim W_i - \dim W_{i-1} - 1$$

for all $i = 0, \ldots, j-1$. Therefore, we have

$$0 \neq \sum_{x \in \Omega} \gamma_S(x b_j w_{j-1} b'_{j-1} \cdots w_1 b'_1 w_0) = |S(b_j w_{j-1} b'_{j-1} \cdots w_1 b'_1 w_0)^{-1}| - |S|. \quad (2)$$

Let U be the isotypic component of V_0 corresponding to an irreducible $\mathbb{Q}G$-module of dimension m. We consider two cases.

First assume that $U \cap W_{j-1} = 0$. Then V_0/W_{j-1} contains an irreducible constituent of dimension m and so $\dim W_{j-1} \leq n-1-m$. Putting $u = b_j w_{j-1} b'_{j-1} \cdots w_1 b'_1 w_0$, we have

$$|u| \leq j + \sum_{i=0}^{j-1} |w_i| \leq j + \sum_{i=0}^{j-1} (\dim W_i - \dim W_{i-1} - 1) = \dim W_{j-1} \leq n - 1 - m.$$

On the other hand, since $P\gamma_S = 0$, it follows that

$$0 = \sum_{x \in \Omega} uP\gamma_S(x) = \sum_{g \in G} (|S(ug)^{-1}| - |S|). \quad (3)$$

Since $|Su^{-1}| - |S| \neq 0$ by equation (2), we conclude by equation (3) that $|S(ug)^{-1}| > |S|$ for some $g \in G$. As u has length at most $n-1-m$ and g can be represented by some word over A of length at most $d_A(G)$, we deduce that $|Sv^{-1}| > |S|$ for some word v over $A \cup B$ of length at most $n - m + d_A(G)$, as required.

Suppose next that $U \cap W_{j-1} \neq 0$ and let $0 \leq k \leq j-1$ be smallest with $U \cap W_k \neq 0$. Then $\dim W_k - \dim W_{k-1} \geq m$. Therefore, putting

$$u' = b_j w_{j-1} b'_{j-1} \cdots w_{k+1} b'_{k+1} \quad \text{and} \quad u'' = b'_k w_{k-1} \cdots w_1 b'_1 w_0,$$

we have that

$$\begin{aligned} |u'u''| &\leq j + \sum_{i \in \{0,\ldots,j-1\} \setminus \{k\}} (\dim W_i - \dim W_{i-1} - 1) \\ &= 1 + \dim W_{j-1} - (\dim W_k - \dim W_{k-1}) \leq n - m. \end{aligned}$$

Using that $Pu''\gamma_S = 0$, as $u''\gamma_S \in W_k \subseteq V_0$, we have that

$$0 = \sum_{x\in\Omega} u'Pu''\gamma_S(x) = \sum_{g\in G}(|S(u'gu'')^{-1}| - |S|). \tag{4}$$

Equation (2) says that $|S(u'w_ku'')^{-1}| - |S| \neq 0$ and hence, as $w_k \in G$, we deduce that $|S(u'gu'')^{-1}| > |S|$ for some $g \in G$ by equation (4). As $|u'u''| \leq n - m$ and g can be represented by some word over A of length at most $d_A(G)$, it follows that $|Sv^{-1}| > |S|$ for some word v over $A \cup B$ of length at most $n - m + d_A(G)$. This completes the proof of the claim.

The theorem now follows, where the final statement is just the observation that $(n-1)^2 = 1 + n(n-2)$. □

Of course, if G is a $\mathbb{Q}$I group, then $m = n - 1$ and so Theorem 9.2 recovers the bound of $1 + (d_A(G) + 1)(n - 2)$ obtained via the spreading property in Theorem 5.16. The weakening of the bound in Theorem 9.2 that replaces m by 1 is essentially contained in the results of Rystsov [89].

As an example, consider S_k acting on the set Ω of 2-sets of $\{1, \ldots, k\}$ with $k \geq 4$. Then $n = |\Omega| = \binom{k}{2}$. It is well known $\mathbb{Q}^\Omega$ is multiplicity-free with three irreducible submodules of dimensions 1, $k-1$ and $\binom{k}{2}-k = n-k$. Theorem 9.2 then shows that if A is any generating set for S_n and $B \subseteq T(\Omega)$ is such that $\langle S_n \cup B\rangle$ is synchronizing (e.g., if k is odd, then any subset B containing a singular map will do), then there is a reset word over $A\cup B$ of length at most $1+(k+d_A(S_k))(n-1)$. In particular, if $d_A(S_k) \leq n - k$, then the Černý bound is achieved. For example, suppose A is the set of Coxeter–Moore generators $(1, 2)$, $(2, 3)$, $\ldots$, $(k - 1, k)$. Then $d_A(S_k) = \binom{k}{2} = n$ and so we obtain a bound of $1 + (k + n)(n - 2) \leq (n-1)^2 + (\sqrt{2n} + 1)(n - 2)$, as $k \leq \sqrt{2n} + 1$.

If s is the number of irreducible constituents (with multiplicities) of V_0, then clearly $ms \geq \dim V_0 = n - 1$. On the other hand, the number of irreducible constituents of the augmentation submodule of the permutation module over $\mathbb{Q}$ is less than the number of irreducible constituents of the augmentation submodule over $\mathbb{C}$ (with multiplicities). If r is the rank of the transitive permutation group G acting on Ω, then $r-1$ is the sum of the squares of the multiplicities of the complex irreducible constituents of the augmentation submodule. Therefore, $r - 1 \geq s$ and so $m \geq \frac{n-1}{r-1}$. Thus Theorem 9.2 admits the following corollary, which avoids representation theoretic language.

Corollary 9.3. *Let G be a transitive permutation group on a set Ω of cardinality $n \geq 2$ and let A be a generating set for G. Suppose that $B \subseteq T(\Omega)$ is such that $\langle G \cup B\rangle$ is synchronizing. Then there is a reset word over $A \cup B$ of length at most*

$$1 + \left(n - \frac{n-1}{r-1} + d_A(G)\right)(n-2)$$

where r is the rank of the permutation group G.

9.2. Regular permutation groups. We next consider regular permutation groups, that is, transitive permutations groups with trivial point stabilizers. Up to isomorphism, this means that we have a finite group G acting on the right of itself, and so for the purpose of this discussion we shall take $\Omega = G$. Notice that the G-invariant graphs in this case are precisely the left Cayley graphs of G with respect to some subset $S \subseteq G$ (not necessarily a generating set). Thus G synchronizes $f \in T(G)$ if and only if f is not an endomorphism of any non-trivial left Cayley graph of G. The only regular permutation groups which are primitive are of prime degree.

In the original paper of Černý [42], the worst case synchronizing automata were constructed by starting with a cyclic permutation of the state set and adjoining an idempotent of rank $n-1$. A cyclic permutation of the states generates a regular permutation group and it is therefore natural to consider in general how quickly regular permutation groups "synchronize" mappings. Here, we are differing from our previous terminology a bit because some of the mappings we adjoin may be permutations, where before we were only adjoining singular mappings. The first result in this subject is due to Rystsov, who proved a slightly more general statement than our formulation [89].

Proposition 9.4 (Rystsov). *Let G be a finite group of order n. Let A be a generating set for G and $B \subseteq T(G)$ such that $\langle G \cup B\rangle$ is synchronizing. Then there is a reset word over $A \cup B$ of length at most $2n^2 - 6n + 5$.*

This can be obtained from Theorem 9.2 by using that $m \geq 1$ and that $d_A(G) \leq n-1$ for a regular permutation group. Proposition 9.4 was improved upon by the third author to the bound in the next theorem, which is sharp in the case of a cyclic group of prime order [97].

Theorem 9.5. *Let G be a finite group of order n and A a generating set for G. Suppose that $B \subseteq T(G)$ is such that $\langle G \cup B\rangle$ is synchronizing. Then there is a reset word over $A \cup B$ of length at most*

$$1 + (n - m(G) + d_A(G))(n-2) \leq 1 + (2n - 1 - m(G))(n-2)$$

where $m(G)$ is the maximum dimension of an irreducible $\mathbb{Q}G$-module.

In particular, in the case that $m(G) \geq d_A(G)$, there is a reset word over $A \cup B$ of length at most $(n-1)^2$.

Theorem 9.5 is immediate from Theorem 9.2 for the case $\Omega = G$ once we make the following observation: the module $V = \mathbb{Q}^G$ is isomorphic to the regular $\mathbb{Q}G$-module and hence contains every irreducible $\mathbb{Q}G$-module as a submodule.

For example, if G is a cyclic group of prime order p, then

$$\mathbb{Q}G \cong \mathbb{Q} \times \mathbb{Q}[x]/(1 + x + \cdots + x^{p-1})$$

and so $m(G) = p - 1$, whereas $d_A(G) \leq p - 1$ for any generating set of G, with equality for a singleton generating set. Thus Theorem 9.5 achieves the Černý bound of $(p-1)^2$ in this case, recovering Pin's theorem [86]. For a cyclic group of

order n, in general, the bound obtained by Theorem 9.5 is $1+(2n-\phi(n)-1)(n-2)$ where ϕ is the Euler totient function.

It is well known that each irreducible representation of the symmetric group S_k over $\mathbb{Q}$ is absolutely irreducible. It follows that S_k has p_k irreducible representations over $\mathbb{Q}$, where p_k is the number of partitions of k, and that the sum of the degrees squared of these representations is $k!$. Thus $p_k m(S_k)^2 \geq k!$ and so we obtain $m(S_k) \geq \sqrt{k!/p_k}$. It is a well-known result of Hardy and Ramanujan that

$$p_k \sim \frac{\exp\left(\pi\sqrt{2k/3}\right)}{4k\sqrt{3}}.$$

On the other hand, Stirling's formula says that $k! \sim \sqrt{2\pi k}\left(\frac{k}{e}\right)^k$. Comparing these expressions, we see that $m(S_k)$ grows faster than any exponential function of k. On the other hand, $d_A(S_k)$ with respect to any of its usual generating sets grows polynomially with k. For instance, if one uses the Coxeter–Moore generators $(1,2)$, $(2,3)$, $\ldots$, $(k-1,k)$ for A, then $d_A(S_k) = \binom{k}{2}$, whereas if one uses the generators $(1,2), (1,2,\ldots,k)$ for A, then $d_A(S_k) \leq (k+1)k(k-1)/2$ since each Coxeter–Moore generator can be expressed as a product of length at most $k+1$ in these generators. Thus Theorem 9.5 yields the Černý bound for either of these generating sets for any k sufficiently large.

Theorem 9.5 was used in [97] to show that if $p \geq 17$ is prime and $B \subseteq T(\mathrm{SL}(2,p))$ is such that $\langle \mathrm{SL}(2,p), B\rangle$ is synchronizing, then there is a reset word over

$$B \cup \left\{ \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} \right\}$$

of length at most $(n-1)^2$ where n is the order of $\mathrm{SL}(2,p)$. Further applications of Theorem 9.5 and its proof idea can be found in [97].

The most elegant result in the Černý conjecture literature is Dubuc's theorem [49].

Theorem 9.6 (Dubuc). *Let Ω be a set of n elements and suppose that $A \subseteq T(\Omega)$ contains a cyclic permutation of Ω. Then if $\langle A\rangle$ is synchronizing, there is a reset word over A of length at most $(n-1)^2$.*

In other words, a cyclic regular permutation group, together with a collection of mappings (some of which may be permutations), synchronizes within the Černý bound whenever it synchronizes, provided that a generator of the cyclic group is one of the input letters for the automaton. Since any generating set of a cyclic group of prime power order must contain a cyclic permutation, we conclude that a cyclic regular permutation group of prime power degree, together with any collection of mappings (not necessarily singular) which generates a synchronizing monoid, always synchronizes within the Černý bound. To formalize this, let us say that a finite group G of order n is a *Černý group* if given any generating set A of G and any subset $B \subseteq T(G)$ such that $A \cup B$ is synchronizing, there is a reset word over $A \cup B$ of length at most $(n-1)^2$.

Theorem 9.6 implies that cyclic groups of prime power order are Černý groups. The third author proved in [97] that elementary abelian p-groups are Černý groups, as are dihedral groups of order $2p$ and $2p^2$ with p an odd prime. Conjecturally, all groups are Černý groups, but this question is far from resolved. Note that since elementary abelian p-groups are Černý groups, it follows that if one takes a synchronizing affine permutation group G, a singular mapping f and a generating set A for G which contains a vector space basis for the subgroup of translations, then there is a reset word for $A \cup \{f\}$ of length at most $(n-1)^2$ where n is the size of the vector space.

As a final comment on Dubuc's theorem, we note the following.

Theorem 9.7. *Let G be a permutation group of non-prime degree n containing an n-cycle. Then G is synchronizing if and only if it is primitive.*

Proof. The forward implication follows from Theorem 3.6; the reverse implication from a theorem of Burnside [113, Theorem 25.3], according to which a primitive group containing a regular cyclic subgroup of composite order is 2-transitive. □

The primitive groups containing a cycle have been classified by Gareth Jones [67].

Note also that there is a growing literature on the diameter of Cayley graphs for certain groups, especially almost simple groups. Babai (see [15]) conjectured that the diameter of any Cayley graph for a simple group G is bounded by a polynomial in $\log |G|$. Such bounds have been found recently for several families of simple groups. However, in these papers the diameter is always in the sense of an undirected graph, whereas we are principally interested in the diameter as a directed graph.

10. Other properties

In this section, we survey briefly another class of permutation groups which lie between the primitive and the 2-homogeneous groups. Whether there is a direct relationship is unknown.

10.1. AS-friendly and AS-free groups. The definition of these classes requires some background on coherent configurations and association schemes. See [3, 36] for more details. The presentation here follows [36].

Coherent configurations were introduced independently by Donald Higman [61, 62] in the USA and by Weisfeiler and Leman [111] in the former Soviet Union to describe the orbits on pairs of a permutation group. Association schemes were introduced earlier by R. C. Bose and collaborators [28, 29] in connection with experimental design in statistics.

Let Ω be a finite set. A *coherent configuration* (c.c.) on Ω is a set $\mathcal{P} = \{R_1, \ldots, R_s\}$ of binary relations on Ω (subsets of Ω^2) satisfying the following four conditions:

(a) $\mathcal{P}$ is a partition of Ω^2;

(b) there is a subset $\mathcal{P}_0$ of $\mathcal{P}$ which is a partition of the diagonal $\Delta = \{(a,a) : a \in \Omega\}$;

(c) for every relation $R_i \in \mathcal{P}$, its *converse* $R_i^\top = \{(b,a) : (a,b) \in R_i\}$ is in $\mathcal{P}$; say $R_i^\top = R_{i^*}$.

(d) there exist integers p_{ij}^k, for $1 \le i,j,k \le s$, such that, for any $(a,b) \in R_k$, the number of points $c \in \Omega$ such that $(a,c) \in R_i$ and $(c,b) \in R_j$ is equal to p_{ij}^k (and, in particular, is independent of the choice of $(a,b) \in R_k$).

A coherent configuration can be defined in terms of its *basis matrices* $A_1, \ldots, A_s$, where A_i is the $\Omega \times \Omega$ matrix with (a,b) entry 1 if $(a,b) \in R_i$, 0 otherwise. In particular, condition (d) asserts that $A_i A_j = \sum_{k=1}^s p_{ij}^k A_k$, so that the span of the basis matrices is an algebra.

If G is any permutation group on Ω, then the partition of Ω^2 into *orbitals* (recall that these are the orbits of G on Ω^2) is a coherent configuration, which we denote by $\mathcal{K}(G)$. We refer to this as the *group case*; a coherent configuration of the form $\mathcal{K}(G)$ is called *Schurian*. In this case, the basis matrices span the centralizer algebra of the permutation representation.

It is clear that a permutation group and its 2-closure define the same coherent configuration, so where necessary we can restrict our attention to 2-closed groups. Indeed, the 2-closure of G is just the automorphism group of $\mathcal{K}(G)$ (the group of permutations preserving all the relations in $\mathcal{K}(G)$).

Let $\mathcal{P}$ be a coherent configuration on Ω. The sets F such that $\{(a,a) : a \in F\}$ belong to $\mathcal{P}$ are called the *fibres* of $\mathcal{P}$; they form a partition of Ω. We say that $\mathcal{P}$ is *homogeneous* if there is only one fibre. If $\mathcal{P} = \mathcal{K}(G)$, the fibres of $\mathcal{P}$ are the orbits of G on Ω; so $\mathcal{K}(G)$ is homogeneous if and only if G is transitive.

A coherent configuration is called *commutative* if its basis matrices commute with one another. It can be shown that, if $\mathcal{P} = \mathcal{K}(G)$, then $\mathcal{P}$ is commutative if and only if the permutation representation is *(complex)-multiplicity-free*.

A coherent configuration is called *symmetric* if all the relations are symmetric. A symmetric c.c. is homogeneous. (For, given any relation R in a c.c. with fibres $F_1, \ldots, F_t$, there are indices i, j such that $R \subseteq F_i \times F_j$.) If $\mathcal{P} = \mathcal{K}(G)$, then $\mathcal{P}$ is symmetric if and only if G is *generously transitive*, that is, any two points of Ω are interchanged by some element of G. Symmetric coherent configurations are also known as *association schemes*, although there is some disagreement over terminology here: each of the four classes of coherent configurations appears with the name "association scheme" somewhere in the literature.

Let $\mathcal{P}$ be a c.c. on Ω. The *symmetrization* $\mathcal{P}^{\mathrm{sym}}$ of $\mathcal{P}$ is the partition of Ω^2 whose parts are all unions of the parts of $\mathcal{P}$ and their converses. It may or may not be a c.c.; if it is, we say that $\mathcal{P}$ is *stratifiable*. The name arises in statistics [16]. It can be shown that, if $\mathcal{P} = \mathcal{K}(G)$, then $\mathcal{P}$ is stratifiable if and only if the permutation representation of G is *real-multiplicity-free*, that is, if it is decomposed into irreducibles over $\mathbb{R}$, they are pairwise non-isomorphic. (Equivalently, the complex irreducibles have multiplicity at most one except for those with Frobenius–Schur index -1, which may have multiplicity 2.)

Thus, the following implications hold:

Proposition 10.1. *A symmetric c.c. is commutative; a commutative c.c. is stratifiable; and a stratifiable c.c. is homogeneous. None of these implications reverses.*

We note also that, if $\mathcal{P} = \mathcal{K}(G)$, then $\mathcal{P}$ is trivial if and only if G is doubly transitive.

To motivate the next definition, we note that the join (in the lattice of partitions) of c.c.s is a c.c.; the same holds for the subclasses defined above. This allows us to define the meet of two c.c.s $\mathcal{C}_1$ and $\mathcal{C}_2$ to be the join of all c.c.s below both of them in the partition lattice; this class is non-empty since the configuration associated with the trivial group (where all parts are singletons) is below any other c.c. However, this does not apply to the subclasses; in particular, there is no meet operation on association schemes.

Let G be a transitive permutation group on the finite set Ω.

(a) We say that G is *AS-free* if the only G-invariant association scheme on Ω is the trivial scheme.

(b) We say that G is *AS-friendly* if there is a unique minimal G-invariant association scheme on Ω.

Of course, if we replaced "AS" by "CC" in the above definitions, then every group would be CC-friendly, and the CC-free groups would be precisely the doubly transitive groups.

Note that a 2-homogeneous group G is AS-free, since the symmetrization of $\mathcal{K}(G)$ is the trivial configuration.

It is easy to see that a uniform partition gives rise to an association scheme (a *group-divisible scheme*), while a Cartesian structure gives rise to an association scheme (a *Hamming scheme*). Thus,

- A transitive permutation group is primitive if and only if it preserves no group-divisible association scheme;
- A primitive permutation group is basic if and only if it preserves no Hamming association scheme.

In a sense, then, the definition of AS-freeness simply carries this idea to its logical conclusion!

Example 10.2. Here is an example of a group which is not AS-friendly. Let G be the symmetric group S_n (for $n \geq 5$), acting on the set Ω of ordered pairs of distinct elements from the set $\{1, \ldots, n\}$: we write the pair (i, j) as ij for brevity. The coherent configuration $\mathcal{K}(G)$ consists of the following parts:

$$\begin{aligned}
R_1 &= \{(ij, ij) : i \neq j\},\\
R_2 &= \{(ij, ji) : i \neq j\},\\
R_3 &= \{(ij, ik) : i, j, k \text{ distinct}\},\\
R_4 &= \{(ij, kj) : i, j, k \text{ distinct}\},\\
R_5 &= \{(ij, ki) : i, j, k \text{ distinct}\},\\
R_6 &= \{(ij, jk) : i, j, k \text{ distinct}\},\\
R_7 &= \{(ij, kl) : i, j, k, l \text{ distinct}\}.
\end{aligned}$$

We have $R_5^\top = R_6$; all other relations are symmetric. The symmetrized partition is not an association scheme, but we find three minimal association schemes as follows:

- the *pair* scheme: $\{R_1, R_2, R_3 \cup R_4, R_5 \cup R_6, R_7\}$;
- two "divisible" schemes $\{R_1, R_3, R_2 \cup R_4 \cup R_5 \cup R_6 \cup R_7\}$ and $\{R_1, R_4, R_2 \cup R_3 \cup R_5 \cup R_6 \cup R_7\}$.

Theorem 10.3. *The following implications hold between properties of a permutation group G:*

$$\begin{array}{ccccccc}
2\text{-}transitive & \Rightarrow & 2\text{-}homogeneous & \Rightarrow & AS\text{-}free & \Rightarrow & primitive\\
\Downarrow & & \Downarrow & & \Downarrow & & \Downarrow\\
gen.\ trans. & \Rightarrow & stratifiable & \Rightarrow & AS\text{-}friendly & \Rightarrow & transitive
\end{array}$$

None of these implications reverses, and no further implications hold.

The smallest 2-closed primitive group which is not AS-friendly is $\mathrm{PSL}(2, 11)$, with degree 55. The smallest 2-closed primitive groups which are AS-friendly but not stratifiable are $\mathrm{PSL}(2, 13)$, with degrees 78 and 91. These groups are numbers $(55, 1)$, $(78, 1)$, $(91, 1)$ and $(91, 3)$ in the list of primitive groups available in GAP. The smallest examples of AS-free groups which are not stratifiable have degree 234, and are isomorphic to $\mathrm{PSL}(3, 3)$ and $\mathrm{PSL}(3, 3) : 2$, numbers $(234, 1)$ and $(234, 2)$ in the list. (Further examples of such groups will be given later.) 2-homogeneous groups which are not generously transitive are well known, as we have seen.

For the class of AS-free groups, we have:

Theorem 10.4. *Let G be a transitive AS-free group. Then G is primitive and basic, and is 2-homogeneous, diagonal or almost simple.*

Almost simple AS-free groups which are not 2-homogeneous do exist. This can be seen from the paper of Faradžev *et al.* [51]. These authors consider the following problem. *Let G be a simple primitive permutation group of order at most 10^6 but not $\mathrm{PSL}(2,q)$. Describe the coherent configurations above $\mathcal{K}(G)$.* Table 3.5.1 on p. 115 gives their results. In several cases, no non-trivial configuration consists entirely of symmetric matrices: such groups are of course AS-free. The smallest example is the group $\mathrm{PSL}(3,3)$, acting on the right cosets of $\mathrm{PO}(3,3)$ (a subgroup isomorphic to S_4), with degree 234; as we have seen, this is the smallest AS-free group which is not 2-homogeneous. Other examples of AS-free groups in this list are the sporadic simple groups M_{12}, degree 1320; J_1, degree 1463, 1540 or 1596; and J_2, degree 1800. The situation is not well understood!

No AS-free primitive diagonal group is known at present. It is known that the socle of such a group must have at least four simple factors. (Groups with two factors preserve a coarsening of the conjugacy class scheme of one factor, while groups with three factors preserve a "Latin square" scheme based on the multiplication table of the factor.)

11. The infinite

We cannot simply take the definition of a synchronizing finite permutation group and extend it to the infinite: there would be no such groups!

Let Ω be an infinite set. Then both the injective maps, and the surjective maps, on Ω form submonoids of the full transformation monoid; they contain the symmetric group but no rank 1 map.

Since the essence of synchronization seems to involve mapping different states to the same place, it is reasonable to require that the map we adjoin is not injective.

Our first attempt at a suitable definition is based on the following fact about the finite case:

Theorem 11.1. *Let M be a transformation monoid on a finite set Ω. Suppose that, for any $v, w \in \Omega$, there exists $f \in M$ with $vf = wf$. Then M is synchronizing.*

Proof. The hypothesis is clearly equivalent to the statement that $\mathrm{Gr}(M)$ is null. (To recall the definition of $\mathrm{Gr}(M)$ see Subsection 4.2.) □

Accordingly, we could try a definition along the following lines:

(a) A transformation monoid M on Ω is *synchronizing* if, for any $v, w \in \Omega$, there exists $f \in M$ with $vf = wf$; equivalently, $\mathrm{Gr}(M)$ is the null graph on Ω.

(b) A permutation group G on Ω is *synchronizing* if, for any map $f\colon \Omega \to \Omega$ which is not injective, the monoid $\langle G, f\rangle$ is synchronizing.

Unfortunately this doesn't give anything interesting!

11.1. Ramsey's Theorem. Ramsey's Theorem is much more general than the form given here; but this is all we need.

Where necessary, we assume the Axiom of Choice, one of whose consequences is that an infinite set contains a countably infinite subset.

Theorem 11.2. *An infinite graph contains either an infinite clique or an infinite independent set.*

By our remark, it suffices to prove this for a countably infinite graph.

Proof. Let $v_1, v_2, \ldots$ be the vertices. We construct inductively a sequence of triples (x_i, Y_i, ϵ_i), where the x_i are distinct vertices, Y_i are infinite decreasing subsets of vertices, $x_i \in Y_j$ if and only if $j < i$, and x_i is joined to all or no vertices of Y_i according as $\epsilon_j = 1$ or $\epsilon_j = 0$. We begin with Y_0 the whole vertex set. Choose $x_i \in Y_{i-1}$. By the Pigeonhole Principle, either x_i has infinitely many neighbours, or it has infinitely many non-neighbours, in Y_{i-1}; let Y_i be the appropriate infinite set and choose ϵ_i appropriately.

Now the sequence $(\epsilon_1, \epsilon_2, \ldots)$ has a constant subsequence; the points x_i corresponding to this subsequence form a clique or independent set, depending on the constant value of ϵ_i. □

We use Ramsey's Theorem to show that the notion of "synchronizing" we just defined is not interesting, at least for permutation groups of countable degree.

Theorem 11.3. *Let G be a permutation group of countable degree. Then G is synchronizing if and only if it is 2-homogeneous.*

Proof. Suppose that G is not 2-homogeneous. Then there is a non-trivial G-invariant graph Γ (take a G-orbit on 2-sets as edges). Replacing Γ by its complement if necessary, and using Ramsey's theorem, we may assume that Γ has a countable clique K. Let v and w be non-adjacent vertices. Choose a bijection f from $\Omega \setminus \{w\}$ to K, and extend it by setting $wf = vf$. Clearly f is an endomorphism of Γ collapsing v and w, and $\langle G, f\rangle$ is not a synchronizing monoid.

Conversely, if G is 2-homogeneous and f a map satisfying $vf = wf$, then $(vg)(g^{-1}f) = (wg)(g^{-1}f)$ for any $g \in G$; so $\langle G, f\rangle$ collapses all pairs, and G is synchronizing. □

11.2. Weak synchronization. We look at a couple of modifications. We say that G is *weakly synchronizing* if, for any map $f\colon \Omega \to \Omega$ of finite rank (that is, having finite image), the monoid $\langle G, f\rangle$ contains a rank 1 map.

Now imprimitive groups may be weakly synchronizing; but it is true that a weakly synchronizing group cannot have a finite system of blocks of imprimitivity. For if S is a transversal for such a system, and f is the map taking any point of Ω to the representative point of f, then $\langle G, f\rangle$ contains no rank 1 map.

Note also that, if M is a transformation monoid containing an element of finite rank, and $\mathrm{Gr}(M)$ is null, then M contains a rank 1 map.

11.3. Strong synchronization. Another possible approach: since, in general, words in $\langle G, f\rangle$ will not be reset words, we should allow infinite words. This requires some preliminary thought.

Let M be a transformation monoid on Ω, and let $\overline{M}$ be its *closure* in the topology of pointwise convergence: a sequence (f_n) of element of M converges to the limit f if, for all $v \in \Omega$, there exists n_0 such that $vf_n = vf$ for all $n \geq n_0$.

Now we say that a permutation group G is *strongly synchronizing* if, for any map f which is not injective, the closure of $M = \langle G, f\rangle$ contains an element of rank 1.

Theorem 11.4. (a) *A strongly synchronizing group is synchronizing.*

(b) *A 2-homogeneous group of countable degree is strongly synchronizing.*

As a consequence of this theorem and the previous one about synchronizing groups, a permutation group of countable degree is strongly synchronizing if and only if it is 2-homogeneous.

Proof. (a) Let f be a map which is not injective, and let (f_n) be a sequence of elements of $\langle G, f\rangle$ converging to a rank 1 function with image $\{z\}$, and choose two distinct points x and y. There exist n_1 and n_2 such that $xf_n = z$ for $n \geq n_1$ and

$yf_n = z$ for $n \geq n_2$. So, if $n = \max(n_1, n_2)$, then $f_n \in \langle G, f\rangle$ and $xf_n = yf_n$. So G is synchronizing.

(b) Let G be 2-homogeneous and let f be a function which is not injective. Choose two points x and y with $xf = yf$. By post-multiplication by an element of G, we can assume that $xf = x$.

Enumerate Ω, as $\{x_1, x_2, \ldots\}$, with $x_1 = x$, and construct a sequence (f_n) of elements of $\langle G, f\rangle$ as follows. Begin with $f_1 = f$. Now suppose that f_n is defined, and satisfies $x_m f_n = x$ for $m \leq n$. If $x_{n+1} f_n = x$, then choose $f_{n+1} = f_n$. Otherwise, choose $g \in G$ mapping $\{x, x_{n+1}\}$ to $\{x, y\}$, and let $f_{n+1} = f_n g f$. Clearly $x_m f_{n+1} = x$ for all $m \leq n+1$. So the sequence converges to the constant function with value x. □

11.4. Larger infinities. Nothing is known about synchronization for larger infinite sets. But the proof that "synchronizing" is equivalent to "2-homogeneous" fails, because of the failure of Ramsey's theorem to guarantee a clique or independent set of the same cardinality as Ω.

We do not know whether the two concepts are equivalent or not for sets of larger cardinalities. The answer might depend on the choice of set-theoretic axioms.

Example 11.5. The Axiom of Choice implies that there is a *well-ordering* of $\mathbb{R}$, a total ordering in which every non-empty subset has a least element. Choose such a well-ordering $\prec$. Now form a graph by joining v and w if $\prec$ and the usual order $<$ agree on $\{v, w\}$, and not if they disagree.

We claim that there is no uncountable clique. Let Y be a clique; then Y is well-ordered by the usual order on $\mathbb{R}$. In a well-order, each non-maximal element v has an immediate successor v'; choose a rational number $q(v)$ in the interval (v, v'). The chosen rationals are all distinct.

Reversing the usual order shows that the complementary graph has the same form; so the graph we constructed has no uncountable independent set either.

11.5. Cores and hulls. The definition of cores in the infinite case is problematic, since it is not clear what "minimal" means. See Bodirsky [27], Bauslaugh [20], for treatments of cores of infinite structures.

Hulls can be defined as usual, but don't do what we want!

Let Γ have vertex set Ω. The *hull* of Γ is the graph $\mathrm{Gr}(\mathrm{End}(\Gamma))$; that is, two vertices v, w are joined in $\mathrm{Hull}(\Gamma)$ if and only if there is no endomorphism f of Γ satisfying $vf = wf$.

Theorem 11.6. *Any countable graph containing an infinite clique is a hull.*

This follows just as in the proof of Theorem 11.3, using Ramsey's theorem.

What happens for graphs with finite clique size? Nick Gravin proved the following result (personal communication from Dima Pasechnik):

Theorem 11.7. *If Γ is an infinite hull with finite clique number, then the clique number and chromatic number of Γ are equal.*

Proof. Let Δ be a finite subset of the vertex set of Γ. If the induced subgraph on Δ is not complete, then there is an endomorphism of Γ collapsing a non-edge of Δ. If Δf_1 is not complete, there is an endomorphism f_2 collapsing a non-edge of Δf_1; and so on. We end up with a homomorphism from A to a complete graph, whose size is at most $\omega(\Gamma)$. So $\chi(\Delta) \leq \omega(\Gamma)$ for every finite subgraph Δ of Γ. It follows from a compactness argument due to de Bruijn (see below) that $\chi(\Gamma) \leq \omega(\Gamma)$. Hence equality holds. □

Theorem 11.8. *Let Γ be an infinite graph, and suppose that every finite subgraph of Γ has chromatic number at most m. Then $\chi(\Gamma) \leq m$.*

Proof. We may suppose Γ countable; let the vertex set be $\{v_1, v_2, \ldots\}$. Construct a graph as follows. Vertices at level i are m-colourings of the induced subgraph on $\{v_1, \ldots, v_i\}$; vertices at levels $i-1$ and i are adjacent if the colouring of $\{v_1, \ldots, v_{i-1}\}$ is a restriction of the colouring of $\{v_1, \ldots, v_i\}$. (The unique vertex at level 0 is the root.) The graph is a tree; each level is finite and non-empty, and there is a path from the root to any vertex. By König's Infinity Lemma, there is an infinite path in the tree, which describes an m-colouring of Γ. □

11.6. Strong primitivity. For infinite groups, Wielandt [112] pointed out a notion which lies between primitivity and 2-transitivity. A permutation group G on Ω is *strongly primitive* if every G-invariant reflexive and transitive relation is trivial (that is, invariant under the symmetric group). Said otherwise, a transitive permutation group G is strongly primitive if and only if every non-diagonal orbital graph for G is strongly connected.

By Theorem 1.1, a finite primitive group is strongly primitive. But, for example, the group $\mathrm{Aut}(\mathbb{Q})$ of order-automorphisms of $\mathbb{Q}$ is primitive (even 2-homogeneous) but not strongly primitive, since the order relation is reflexive and transitive but not symmetric.

We can refine Wielandt's notion by saying that a permutation group G on Ω is *strong* if every G-invariant reflexive and transitive relation is symmetric (and so is an equivalence relation). For example, a *torsion group* (one in which all elements have finite order) is strong.

Now G is strongly primitive if and only if it is strong and primitive.

12. Problems

In this section we propose a number of problems that are naturally prompted by the results in this paper.

The next problem might be appropriate for a PhD project.

Problem 12.1. *Find the synchronizing affine groups of degrees p^2 or p^3, with p prime.*

The finite 2-transitive and 2-homogeneous groups are known; $\mathbb{Q}$I groups are also known. What we do not know is whether there are spreading non-$\mathbb{Q}$I groups.

Problem 12.2. *Is there any group which is spreading but not* $\mathbb{Q}$I*?*

If the previous question turns out to have a negative answer, then the next two problems will have a trivial answer.

Problem 12.3. *Is there an infinite family of groups which are spreading but not* $\mathbb{Q}$I*?*

Problem 12.4. *Classify the spreading groups.*

It is also natural to consider the class of strongly separating groups. Call a transitive permutation group G on Ω *strongly separating* if whenever A, B are non-trivial subsets of Ω such that $|B|$ divides $|\Omega|$ and $|A|\cdot|B| = k\cdot|\Omega|$, then there exists $g \in G$ such that $|A \cap Bg| \neq k$. Clearly, spreading groups are strongly separating and strongly separating groups are separating. Note that S_m acting on 2-sets is separating for m odd, but not strongly separating, and that an affine group is strongly separating if and only if it is $\mathbb{Q}$I. Also, S_m acting on k-sets with $k \geq 3$ is never strongly separating and if the Hadamard conjecture is true, then S_{2m} acting on (m, m) partitions is not strongly separating.

One can show that G is strongly separating if and only if, for each singular mapping f with uniform kernel, and each proper subset B of Ω that is a union of $\mathrm{Ker}(f)$-classes, there exists $g \in G$ such that $|Bgt^{-1}| > |B|$. It follows that if G is spreading and f is a uniform singular mapping with $|\Omega f| = d$, then there is a reset word over $G \cup \{f\}$ with at most d occurrences of f.

Problem 12.5. *Are there strongly separating groups that are not spreading? Are there strongly separating groups that are not* $\mathbb{Q}$I*?*

We know that separating groups properly contain the class of spreading groups.

Problem 12.6. *Classify separating groups modulo a classification of spreading groups.*

Similarly we know that synchronizing groups properly contain separating groups. However we only have one infinite family (Example 6.9) and one sporadic example.

Problem 12.7. *Is there another infinite family of groups which are synchronizing but not separating?*

A particular instance of the previous problem is the following.

Problem 12.8. *Is it true that* $\mathrm{P}\Omega(5, q)$*, for q odd, form an infinite family of synchronizing, but not separating groups?*

This is equivalent to asking for a proof that the polar spaces associated with these groups do not have a partition into ovoids. As in the example, this would follow if it could be shown that ovoids in this space are necessarily classical (that is, hyperplane sections); but it could hold even if non-classical ovoids exist.

Of course, in this respect, the ultimate goal would be an answer for the following problem.

Problem 12.9. *Classify synchronizing groups modulo a classification of separating groups.*

A particular instance of the previous problem is the following.

Problem 12.10. *Is it true that in the class of affine groups, synchronizing and separating coincide?*

Problem 12.11. *Do there exist subsets A and B of a finite simple group S, neither of which is a coset of a subgroup, such that $|g^{-1}Ah \cap B| = 1$ for all $g, h \in S$?*

Since the symmetric group S_m acting on pairs of points of $\{1, \ldots, m\}$, for m even, is basic and almost synchronizing, but not synchronizing, it follows that the intersection of the former two classes properly contain the latter. More examples of almost synchronizing, but not synchronizing groups, can be found in [3].

Problem 12.12. *Classify basic almost synchronizing groups modulo a classification of synchronizing groups.*

We already saw that there are basic not almost synchronizing groups, and there are non-basic almost synchronizing groups.

Problem 12.13. *Classify basic groups modulo a classification of the basic and almost synchronizing groups.*

Problem 12.14. *Classify almost synchronizing groups modulo a classification of the basic and almost synchronizing groups.*

A slight variation of the previous problems gives the following.

Problem 12.15. *Classify almost synchronizing groups that are not basic.*
Classify basic groups that are not almost synchronizing.

A first step would be to decide whether the wreath product S_k wr S_m (in the power action) is almost synchronizing. It is known that there are uniform maps of rank k^i for $i = 1, 2, \ldots, m-1$ which are not synchronized by this group; but does it synchronize any non-uniform maps?

Since primitive groups properly contain basic and almost synchronizing groups, the following problem is natural.

Problem 12.16. *Classify primitive groups modulo the classification of basic groups and almost synchronizing groups.*

Our last questions on the relations of these groups are the following.

Problem 12.17. *How does almost synchronizing relate to partition separating?*

Problem 12.18. *How does basic relate to partition separating?*

Problem 12.19. *Classify partition separating modulo a classification of almost synchronizing and basic.*

Many of the above classification problems will be difficult, since they include notorious unsolved problems in extremal combinatorics, design theory and finite geometry.

Problem 12.20. *Is there a sublinear bound for the number of non-synchronizing ranks of a primitive group?*

A weaker (and perhaps easier) question would be:

Problem 12.21. *Is there a sublinear bound for the number of ranks of endomorphisms of a vertex-primitive graph?*

A graph Γ is a *pseudocore* if every endomorphism is either an automorphism or a proper colouring. Clearly the automorphism group of a pseudocore (if it is transitive) is almost synchronizing and has only one non-synchronizing rank. Several examples of such graphs are known [57]. Indeed, there is no known graph with primitive automorphism group with permutation rank 3 which is not a pseudocore. If it is true that every strongly regular graph is a pseudocore, as a recent preprint by David Roberson claims (personal communication from Gordon Royle), it would follow that if G is primitive with permutation rank 3, then $|\mathrm{NS}(G)| \leq 1$. (Recall that the permutation rank is the number of G orbits on Ω^2.)

Problem 12.22. *Is it true that, for any primitive permutation group G, $|\mathrm{NS}(G)|$ is bounded by a function of the permutation rank?*

Regarding closures, we ask the following.

Problem 12.23. *Is it true that the $\mathbb{R}$I-closure of a permutation group is equal to its strong 2-closure?*

Problem 12.24. *What properties does the $\mathbb{F}$I-closure of a permutation group has when $\mathbb{F}$ is a field of non-zero characteristic, or a local field such as the p-adic numbers?*

The following question, related to closures, has been nagging the third author for years.

Problem 12.25. *Let $M \subseteq T(\Omega)$ be a transformation monoid. What is the relationship between 2-transitivity of M and irreducibility of the augmentation submodule of $\mathbb{C}\Omega$? Note that if M is not a group, then both properties imply that M is primitive and synchronizing. (A transformation monoid is* primitive *if there are no equivalence relations $\equiv$ on Ω such that $x \equiv y$ implies $xm \equiv ym$ for all $m \in M$ except for the equality and universal relations.)*

Regarding the groups linked to association schemes, we have a number of problems that parallel those above about synchronizing groups.

Problem 12.26. *Is there any AS-free primitive diagonal group?*

More generally, we have the following.

Problem 12.27. *Classify AS-free groups.*

Problem 12.28. *Classify stratifiable groups modulo a classification of generously transitive groups.*

Problem 12.29. *Classify AS-friendly groups modulo a classification of AS-free and stratifiable groups.*

We introduced two hierarchies in this paper: one on page 56, involving synchronizing groups and friends, and another in Theorem 10.3.

Problem 12.30. *Draw a single* Venn Diagram *of the two hierarchies, together with partition separating groups and strongly separating groups, pointing out which regions of the diagram are not empty, contain infinite families, etc.*

There are efficient algorithms to check if a given set of permutations generates a primitive group.

Problem 12.31. *Find an efficient [polynomial-time] algorithm to decide if a given set of permutations generates a synchronizing [spreading, separating, almost synchronizing, partition separating, AS-free, generously transitive, stratifiable, AS-friendly] group.*

Existing algorithms for deciding whether a given primitive group is synchronizing or separating involve solving NP-hard problems, such as clique number and chromatic number, but on rather special graphs (those which are vertex-primitive). Does the information available about primitive groups using CFSG allow faster algorithms to be found?

Problem 12.32. *What is the computational complexity of computing the* 2*-closure of a (primitive) permutation group?*

The previous problems are linked to the following problem.

Problem 12.33. *For the extent of GAP's library of primitive groups, include in GAP the list of synchronizing [spreading, separating, almost synchronizing, AS-free, generously transitive, stratifiable, AS-friendly] groups.*

The two first named authors, Gordon Royle, James Mitchell [82], Artur Schafer, Csaba Schneider, Leonard Soicher, and Pablo Spiga, independently, wrote GAP code that produced lists for some of these classes (almost synchronizing, synchronizing, spreading, separating, or the association schemes classes), reaching, in the best cases (Royle's and Spiga's), degrees of a few hundreds. In order to better understand all these classes of groups, examples of larger degrees are needed and more sophisticated code or algorithms must come into play.

Recently, Leonard Soicher (private communication to the two first named authors) produced GAP code that finds non-synchronizing groups one order of magnitude faster than any other previous tool (about 20 seconds to find all non-synchronizing groups up to degree 100).

Problem 12.34. *For the extent of GAP's library of primitive groups, include in GAP a library of all pairs of partitions and sections preserved by a given non-synchronizing group.*

Problem 12.35. *For the extent of GAP's library of primitive groups, include in GAP a library of all classes of groups in this paper (almost synchronizing, synchronizing, spreading, separating, strongly separating, the association schemes classes, etc).*

There are a number of natural problems relating the subject of this text to the Černý conjecture.

Problem 12.36. *Is it true that if $A \subseteq S_n$ generates a primitive group and f is a rank $n-1$ mapping, then there is a reset word over $A \cup \{f\}$ of length at most $(n-1)^2$? Rystsov proved a quadratic bound if f is an idempotent, but John Dixon (in an unpublished example) showed that Rystsov's proof scheme cannot yield the bound of $(n-1)^2$, even in the case that A generates the symmetric group for $n \geq 5$!*

Problem 12.37. *Is every group a Černý group? Perhaps one can generalize Dubuc's scheme* [49] *to abelian groups; however, it is not immediately clear how to generalize Dubuc's result to arbitrary generating sets of cyclic groups of order not a prime power!*

Problem 12.38. *Is the Černý conjecture true for synchronizing automata such that some subset of the transitions generates a transitive permutation group? How about the same question replacing the adjective "transitive" by any of the following adjectives: primitive, synchronizing, separating, spreading, $\mathbb{Q}$I, 2-homogeneous or 2-transitive?*

Problem 12.39. *Let G be a primitive group contained in S_n, and generated by $S \subseteq G$. Let P be a partition of $\{1, \ldots, n\}$. Is it true that if G takes a 2-subset of X into some part of P, then it can do so in a word over S of size linear in n?*

The best general bound so far for the length of a reset word in a synchronizing automaton was proved by Pin [86] and Frankl [54]; see also [74]. The idea is the following. Suppose we have a reset word w of minimal length. We have an n-set X and $|Xw| = 1$. Let k be arbitrary in $1 < k < n$. Then there are two prefixes of w, say w_k and w'_k, such that $|Xw_k| \geq k$ and $|Xw_{k+1}| < k$, and w_k, w'_k are the smallest prefixes satisfying these properties. In particular, $w'_k = w_k a_1 \ldots a_m$. Let $P_i := Xw_k a_1 \ldots a_i$. Clearly $|P_i| = k$, for all $i \in \{1, \ldots, m-1\}$. Since $|P_{m-1}| \geq k$ and $|P_m| < k$, there are two elements $p, q \in P_{m-1}$ in the same kernel class of a_m. Let $x^i := x a_{m-1}^{-1} \ldots a_i^{-1}$, for $x \in \{p, q\}$ and $i \in \{2, \ldots, m-1\}$. Observe that $x^i \in P_{i-1}$, and, given the minimality of w, $i-1$ is the smallest index of the Ps containing both p^i and q^i. Therefore Pin asked how large can be m in a family of sets subject to this condition; the answer to this question by Frankl [54] yielded the current best bound for reset words (please see also [108]).

However, if we are dealing with a spreading group in which the non-invertible map only comes into play to reduce the rank, all the sets P_i (for $i \in \{1, \ldots, m-1\}$), in the argument above, are in the orbit of a k-set. Therefore, hopefully, the bound will be much lower than the one found by Frankl [54]. Therefore the following problem is very natural.

Problem 12.40. *Solve the problem analogous to the one solved in* [54] *but with the extra hypothesis that the k-sets are all contained in the same orbit under the action of some spreading group.*

A final general problem on the philosophy behind our investigation in this paper:

Problem 12.41. *We have seen in this paper several instances where a characterization of primitivity leads to a condition which can be generalized leading to a new class of permutation groups. For example,*

- *G is primitive if and only if it synchronizes every map of rank $n-1$; this leads us to the class of synchronizing groups.*
- *G is primitive if and only if it preserves no divisible association scheme; this leads us to the class of AS-free groups.*
- *G is primitive if and only if every orbital graph for G is connected; replacing "connected" by "strongly connected" leads to the class of strongly primitive groups.*
- *G is primitive if and only if, for every 2-set A and 2-partition P, there exists $g \in G$ such that Ag is a section for P. (This is the assertion that a graph is connected if, for every 2-partition, there is an edge of the graph which is a section.) If we replace 2 by k here, we get the definition of the* k-universal transversal property*, which is discussed (together with its implications for semigroup theory) in* [13].

In the authors' view, it is worthwhile searching for other characterizations of primitivity, in the hope of uncovering other interesting classes to study!

Acknowledgement

The authors thank the referee for a very thorough and thoughtful report, which has materially improved the clarity and accuracy of this paper.

References

[1] R. L. Adler and B. Weiss, Similarity of automorphisms of the torus, *Memoirs of the American Mathematical Society* **98**, American Mathematical Society, Providence, R.I., 1970.

[2] R. L. Adler, W. L. Goodwyn and B. Weiss, Equivalence of topological Markov shifts, *Israel J. Math.* **27** (1977), 48–63.

[3] P. P. Alejandro, R. A. Bailey and P. J. Cameron, Association schemes and permutation groups, *Discrete Math.* **266** (2003), 47–67.

[4] J. Almeida and B. Steinberg, Matrix mortality and the Černý-Pin conjecture, In *Developments in language theory*, volume 5583 of *Lecture Notes in Comput. Sci.*, pages 67–80. Springer, Berlin, 2009.

[5] Jorge Almeida, Stuart Margolis, Benjamin Steinberg, and Mikhail Volkov, Representation theory of finite semigroups, semigroup radicals and formal language theory. *Trans. Amer. Math. Soc.*, **361** (2009), 1429–1461.

[6] D. S. Ananichev and M. V. Volkov, Some results on Černý type problems for transformation semigroups. In *Semigroups and languages*, pages 23–42. World Sci. Publ., River Edge, NJ, 2004.

[7] D. S. Ananichev and M. V. Volkov, Synchronizing generalized monotonic automata. *Theoret. Comput. Sci.*, **330** (2005), 3–13.

[8] D. S. Ananichev, M. V. Volkov, and V. V. Gusev, Primitive digraphs with large exponents and slowly synchronizing automata. *Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI)*, 402(Kombinatorika i Teoriya Grafov. IV):9–39, 218, 2012.

[9] D. S. Ananichev, M. V. Volkov, and Yu. I. Zaks, Synchronizing automata with a letter of deficiency 2. *Theoret. Comput. Sci.*, **376** (2007), 30–41.

[10] J. Araújo, A group theoretical approach to synchronizing automata and the Černy problem. Unpublished manuscript, 2006.

[11] J. Araújo, W. Bentz, P. J. Cameron, G. Royle and A. Schaefer, Primitive groups and synchronization, `http://arxiv.org/abs/1504.01629`

[12] J. Araújo and P. J. Cameron, Primitive groups synchronize non-uniform maps of extreme ranks, *Journal of Combinatorial Theory, Series B*, **106** (2014), 98–114.

[13] J. Araújo and P. J. Cameron, Two generalizations of homogeneity in groups with applications to regular semigroups, *Trans. Amer. Math. Soc.*, in press (published on-line 1 July 2015).

[14] F. Arnold and B. Steinberg, Synchronizing groups and automata, *Theoretical Computer Science* **359** (2006), 101–110.

[15] László Babai and Ákos Seress, On the diameter of permutation groups, *European J. Combinatorics* **13** (1992), 231–243.

[16] R. A. Bailey, Strata for randomized experiments, *J. Royal Statist. Soc.* (B) **53** (1991), 27–78.

[17] S. Ball, P. Govaerts and L. Storme, On ovoids of parabolic quadrics, *Designs, Codes, Cryptography* **38** (2006), 131–145.

[18] J. Bamberg, M. Giudici, M. W. Liebeck, C. E. Praeger and J. Saxl, The classification of almost simple $\frac{3}{2}$-transitive groups, http://arxiv.org/abs/1103.6069

[19] Zs. Baranyai, On the factorization of the complete uniform hypergraph, *Infinite and finite sets* (Colloq., Keszthely, 1973), Vol. I, pp. 91–108, *Colloq. Math. Soc. Janos Bolyai*, **10**, North-Holland, Amsterdam, 1975.

[20] Benjamin Bauslaugh, Core-like properties of infinite graphs and structures, *Discrete Math.* **138** (1995), 101–111.

[21] Marie-Pierre Béal, Mikhail V. Berlinkov, and Dominique Perrin, A quadratic upper bound on the size of a synchronizing word in one-cluster automata. *Internat. J. Found. Comput. Sci.*, **22** (2011), 277–288.

[22] J. L. Berggren, An algebraic characterization of finite symmetric tournaments, *Bull. Austral. Math. Soc.* **6** (1972), 53–59.

[23] M. V. Berlinkov. On the probability of being synchronizable. http://arxiv.org/abs/1304.5774

[24] A. Beutelspacher, On parallelisms in finite projective spaces, *Geometriae Dedicata* **3** (1974), 35–40.

[25] N. L. Biggs, Three remarkable graphs, *Canad. J. Math.* **25** (1973), 397–411.

[26] B. J. Birch, R. G. Burns, S. O. Macdonald and P. M. Neumann, On the orbit-sizes of permutation groups containing elements separating finite subsets, *Bull. Austral. Math. Soc.* **14** (1976), 7–10.

[27] Manuel Bodirsky, The cores of a countably categorical structure, In: Diekert V., Durand B. (eds) STACS 2005, *Lecture Notes in Computer Science*, vol 3404. Springer, Berlin, Heidelberg, 2005.

[28] R. C. Bose and D. M. Mesner, On linear associative algebras corresponding to association schemes of partially balanced designs, *Ann. Math. Statist.* **30** (1959), 21–38.

[29] R. C. Bose and T. Shimamoto, Classification and analysis of partially balanced designs with two associate classes, *J. Amer. Statist. Assoc.* **47** (1952), 151–184.

[30] Michael Braun, Tuvi Etzion, Patric Østergård, Alexander Vardy and Alfred Wasserman, Existence of q-analogs of Steiner systems, http://arxiv.org/abs/1304.1462

[31] W. Burnside, On some properties of groups of odd order, *Proc. London Math. Soc.* **33** (1901), 162–185.

[32] Qian Cai and Hua Zhang, A note on primitive permutation groups of prime power degree, *J. Discrete Math.* #194741, http://dx.doi.org/10.1155/2015/194741

[33] P. J. Cameron, Bounding the rank of certain permutation groups, *Math. Z.* **124** (1972), 343–352.

[34] P. J. Cameron, *Projective and Polar Spaces*, QMW Maths Notes 13, QMW, London, 1991; available from https://cameroncounts.files.wordpress.com/2015/04/pps1.pdf

[35] P. J. Cameron, *Permutation Groups*, London Math. Soc. Student Texts **45**, Cambridge University Press, Cambridge, 1999.

[36] P. J. Cameron, Coherent configurations, association schemes and permutation groups, pp. 55–71 in Groups, Combinatorics and Geometry (ed. A. A. Ivanov, M. W. Liebeck and J. Saxl), World Scientific, Singapore, 2003.

[37] P. J. Cameron, Dixon's Theorem and random synchronization, *Discrete Mathematics* **313** (2013), 1233–1236.

[38] P. J. Cameron and P. A. Kazanidis, Cores of symmetric graphs, *J. Austral. Math. Soc.* **85** (2008), 145–154.

[39] P. J. Cameron, P. M. Neumann and D. N. Teague, On the degrees of primitive permutation groups, *Math. Z.* **180** (1982), 141–149.

[40] A. Carpi and F. d'Alessandro, The synchronization problem for strongly transitive automata. In *Developments in language theory*, volume 5257 of *Lecture Notes in Comput. Sci.*, pages 240–251. Springer, Berlin, 2008.

[41] A. Carpi and F. d'Alessandro, The synchronization problem for locally strongly transitive automata. In *Mathematical Foundations of Computer Science*, volume 5734 of *Lecture Notes in Comput. Sci.*, pages 211–222. Springer, Berlin, 2009.

[42] J. Černý, Poznámka k homogénnym eksperimentom s konecnými automatami [A remark on homogeneous experiments with finite automata], *Mat.-Fyz. Časopis Slovensk. Akad. Vied.* **14** (1964), 208–216.

[43] C. J. Colbourn and L. Zhu, The spectrum of R-orthogonal Latin squares, *Combinatorics advances*, Math. Appl, **329**, 1995, 49–75.

[44] H. S. M. Coxeter, The chords of the non-ruled quadric in PG$(3, 3)$, *Canad. J. Math.* **10** (1958), 484–488.

[45] J. De Beule, Substructures of finite classical polar spaces. In Current research topics in Galois geometry, Mathematics Research Developments, chapter 2, pages 35-61. NOVA Sci. Publ., New York, 2012. `http://homepages.vub.ac.be/~jdbeule/postprints/DBKM_NOVA2012.pdf`

[46] J. D. Dixon, Permutations representations and rational irreducibility, *Bull. Austral. Math. Soc.* **71** (2005), 493–503.

[47] J. D. Dixon and B. Mortimer, *Permutation Groups*, Springer, 1996.

[48] R. G. Downey, and M. R. Fellows, Fixed-parameter tractability and completeness, II: On completeness for W[1], *Theoretical Computer Science* **141** (1995), 109–131.

[49] L. Dubuc, Sur les automates circulaires et la conjecture de Černý. *RAIRO Inform. Théor. Appl.* **32** (1998), 21–34.

[50] David Eppstein, Reset sequences for monotonic automata, *SIAM J. Comput.* **19** (1990), 500–510.

[51] I. A. Faradžev, M. H. Klin and M. E. Muzichuk, Cellular rings and groups of automorphisms of graphs, in *Investigations in Algebraic Theory of Combinatorial Objects* (I. A. Faradžev, A. A. Ivanov, M. H. Klin and A. J. Woldar, eds.), Kluwer, Dordrecht, 1994, pp. 1–152.

[52] J. S. Frame, The degrees of the irreducible components of simply transitive permutation groups, *Duke Math. J.* **3** (1937), 8–17.

[53] J. S. Frame, The double cosets of a finite group, *Bull. Amer. Math. Soc.* **47** (1941), 458–467.

[54] P. Frankl, An extremal problem for two families of sets, *European Journal of Combinatorics* **3** (1982), 125–127.

[55] Ian P. Gent, Chris Jefferson and Ian Miguel, MINION: A Fast, Scalable, Constraint Solver, (slides) in Proceedings of the 17th European Conference on Artificial Intelligence (ECAI 2006); available from `https://blogs.cs.st-andrews.ac.uk/constraintmodelling/files/2015/07/MinionECAI06.pdf`

[56] The GAP Group, GAP – Groups, Algorithms, and Programming, Version 4.8.6; 2016. `http://www.gap-system.org`

[57] C. D. Godsil and G. F. Royle, Cores of geometric graphs, *Ann. Combinatorics* **15** (2011), 267–276.

[58] G. Hahn and C. Tardif, Graph homomorphisms: structure and symmetry pp. 107–166 in *Graph Symmetry: Algebraic Methods and Applications* (ed. G. Hahn and G. Sabidussi), Kluwer, 1997.

[59] P. Hell and J. Nešetřil, *Graphs and Homomorphisms*, Oxford University Press, Oxford, 2004.

[60] Christoph Hering, Transitive linear groups and linear groups which contain irreducible subgroups of prime order, *Geometriae Dedicata* **2** (1974), 425–460.

[61] D. G. Higman, Intersection matrices for finite permutation groups, *J. Algebra* **6** (1967), 22–42.

[62] D. G. Higman, *Combinatorial Considerations about Permutation Groups*, Mathematical Institute, Oxford, 1971.

[63] J. W. P. Hirschfeld and J. A. Thas, *General Galois Geometries*, 2nd edition, Springer Monographs in Mathematics, Springer-Verlag, London, 2016.

[64] J.M. Howie, The subsemigroup generated by the idempotents of a full transformation semigroup, *J. London Math. Soc.*, **41**, (1966), 707–716.

[65] Bertram Huppert, Zweifach transitive, auflösbare Permutationsgruppen, *Math. Z.* **68** (1957), 126–150.

[66] K. W. Johnson, S-rings over loops, right mapping groups and transversals in permutation groups, *Math. Proc. Cambridge Philos. Soc.* **89** (1981), 433–443.

[67] Gareth A. Jones, Primitive permutation groups containing a cycle, *Bull. Austral. Math. Soc.* **89** (2014), 159–165.

[68] Raphaël M. Jungers, The synchronizing probability function of an automaton. *SIAM J. Discrete Math.* **26** (2012), 177–192.

[69] W. M. Kantor, Automorphism groups of designs, *Math. Z.* **109** (1969), 246–252.

[70] Jarkko Kari, A counter example to a conjecture concerning synchronizing words in finite automata. *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS* **73** (2001), 146.

[71] Jarkko Kari, Synchronizing finite automata on Eulerian digraphs. *Theoret. Comput. Sci.* **295** (2003), 223–232. Mathematical foundations of computer science (Mariánské Lázně, 2001).

[72] P. Keevash, The existence of designs, `http://arxiv.org/abs/1401.3665`

[73] H. Kharaghani and B. Tayfeh-Rezaie, A Hadamard matrix of order 428, *J. Combinatorial Designs* **13** (2005), 435–440.

[74] A. A. Klyachko, I.C. Rystsov, and M. A. Spivak, On an extremal combinatorial problem connected with an estimate for the length of a reflexive word in an automaton, *Kibernetika (Kiev)* **2** (1987), 16–20, 25, 132.

[75] E. A. Kuznetsov, Transversals in groups, I, Elementary properties, *Quasigroups and related systems* **1** (1994), 22–42.

[76] Cai Heng Li and Ákos Seress, The primitive permutation groups of squarefree degree, *Bull. London Math. Soc.* **35** (2003), 635–44.

[77] Martin W. Liebeck, The affine permutation groups of rank three, *Proc. London Math. Soc.* (3) **54** (1987), 477–516.

[78] M. W. Liebeck, C. E. Praeger and J. Saxl, On the O'Nan-Scott theorem for finite primitive permutation groups, *J. Austral. Math. Soc. Ser. A* **44** (1988), 389–396.

[79] M. W. Liebeck, C. E. Praeger and J. Saxl, The maximal subgroups of the finite simple groups and their automorphism groups, *Memoirs Amer. Math. Soc.* **86** (1990), 1–151.

[80] L. Lovász, Kneser's conjecture, chromatic number, and homotopy, *J. Combinatorial Theory* (A) **25** (1978), 319–324.

[81] Donald B. McAlister, Semigroups generated by a group and an idempotent, *Comm. Algebra* **26** (1998), 515–547.

[82] J. D. Mitchell and others, Semigroups - GAP package, Version 2.6, (2015). http://www-groups.mcs.st-andrews.ac.uk/ jamesm/semigroups.php

[83] Peter M. Neumann, The lawlessness of groups of finitary permutations. *Arch. Math.* (Basel) **26** (1975), 561–566.

[84] Peter M. Neumann, Primitive permutation groups and their section-regular partitions, *Michigan Math. J.* **58** (2009), 309–322.

[85] C. Nicaud, Fast Synchronization of Random Automata. `http://arxiv.org/abs/1404.6962`

[86] J.-E. Pin, Sur un cas particulier de la conjecture de Cerny. In *Automata, languages and programming (Fifth Internat. Colloq., Udine, 1978)*, volume 62 of *Lecture Notes in Comput. Sci.*, pages 345–352. Springer, Berlin, 1978.

[87] J.-E. Pin, On two combinatorial problems arising from automata theory. In *Combinatorial mathematics (Marseille-Luminy, 1981)*, volume 75 of *North-Holland Math. Stud.*, pages 535–548. North-Holland, Amsterdam, 1983.

[88] I. C. Rystsov, On the rank of a finite automaton. *Kibernet. Sistem. Anal.*, **3** (1992), 3–10, 187.

[89] I. K. Rystsov, Quasioptimal bound for the length of reset words for regular automata. *Acta Cybernet.* **12** (1995), 145–152.

[90] Igor Rystsov, Reset words for commutative and solvable automata. *Theoret. Comput. Sci.* **172** (1997), 273–279.

[91] I. K. Rystsov, Estimation of the length of reset words for automata with simple idempotents, *Cybernetics and Systems Analysis* **36** (2000), 339–344.

[92] Arto Salomaa, Composition sequences for functions over a finite domain. *Theoret. Comput. Sci.* **292**, 263–281. Selected papers in honor of Jean Berstel.

[93] C. Schneider and A. C. Silva, Cliques and colorings in generalized Paley graphs and an approach to synchronization. *J. Algebra Appl.* **14** (2015), no. 6.

[94] I. Schur, Über die Kongruenz $x^m + y^m \equiv z^m \pmod{p}$, *Jahresber. Deutsch Math.-Verin.* **25** (1916), 114–116.

[95] J.-P. Serre, On a theorem of Jordan, *Bull. Amer. Math. Soc.* **40** (2003), 429–440.

[96] Richard P. Stanley, *Enumerative Combinatorics*, Vol. 2. Cambridge Studies in Advanced Mathematics **62**, Cambridge University Press, Cambridge, 1999.

[97] Benjamin Steinberg, Černý's conjecture and group representation theory, *J. Algebraic Combinatorics* **31** (2010), 83–109.

[98] Benjamin Steinberg, A theory of transformation monoids: combinatorics and representation theory, *Electron. J. Combin.* **17** (2010), Research Paper 164, 56 pp. (electronic).

[99] Benjamin Steinberg, The averaging trick and the Černý conjecture. *Internat. J. Found. Comput. Sci.* **22** (2011), 1697–1706.

[100] Benjamin Steinberg, The Černý conjecture for one-cluster automata with prime length cycle. *Theoret. Comput. Sci.* **412** (2011), 5487–5491.

[101] Benjamin Steinberg, *The Representation Theory of Finite Monoids*, Springer, to appear.

[102] D. E. Taylor, *The Geometry of the Classical Groups*, Helderman, Berlin, 1992.

[103] L. Teirlinck, A completion of Lu's determination of the spectrum for large sets of disjoint Steiner triple systems. *J. Combinatorial Theory* (A) **57** (1991), 302–305.

[104] A. N. Trahtman, An efficient algorithm finds noticeable trends and examples concerning the Černy conjecture. In *Mathematical foundations of computer science 2006*, volume 4162 of *Lecture Notes in Comput. Sci.*, pages 789–800. Springer, Berlin, 2006.

[105] A. N. Trahtman, The Černý conjecture for aperiodic automata. *Discrete Math. Theor. Comput. Sci.* **9** (2007), 3–10 (electronic).

[106] Avraham Trahtman, The road coloring problem, *Israel J. Math.* **172** (2009), 51–60.

[107] W. T. Tutte, The chords of the non-ruled quadric in PG(3, 3), *Canad. J. Math.* **10** (1958), 481–483.

[108] M. V. Volkov, Synchronizing automata and the Černý conjecture, LATA 2008, LNCS 5196 (2008), 11–27.

[109] S. Wagstaff, Divisors of Mersenne numbers, *Math. Comp.* **40** (1983), 385–397.

[110] W. D. Wallis, Anne Penfold Street, and Jennifer Seberry Wallis, *Combinatorics: Room Squares, Sum-Free Sets, Hadamard Matrices*, Lecture Notes in Math. **292**, Springer, Berlin, 1972.

[111] B. Yu. Weisfeiler and A. A. Leman, Reduction of a graph to a canonical form and an algebra which appears in the process, *Scientific-Technological Investigations* (2) **9** (1968), 12–16.

[112] H. Wielandt, *Unendliche Permutationsgruppen*, Lecture Notes, Universität Tübingen, 1959. [English translation by P. Bruyns included in Wielandt's collected works: H. Wielandt, *Mathematische Werke: Mathematical Works*, Volume 1 (Bertram Huppert and Hans Schneider, eds.), Walter de Gruyter, Berlin, 1994, pp. 199–235.]

[113] H. Wielandt, *Finite Permutation Groups*, Academic Press, New York, 1964. [Included in Wielandt's collected works: H. Wielandt, *Mathematische Werke: Mathematical Works*, Volume 1 (Bertram Huppert and Hans Schneider, eds.), Walter de Gruyter, Berlin, 1994, pp. 119–198.]

[114] H. Wielandt, *Permutation groups through invariant relations and invariant functions*, Lecture Notes, Ohio State University, 1969. [Included in Wielandt's collected works: H. Wielandt, *Mathematische Werke: Mathematical Works*, Volume 1 (Bertram Huppert and Hans Schneider, eds.), Walter de Gruyter, Berlin, 1994, pp. 237–296.]

[115] R. A. Wilson, *The Finite Simple Groups*, Graduate Texts in Mathematics 251, Springer, 2009.

[116] L. Zhu and H. Zhang, Completing the spectrum of r-orthogonal Latin squares, *Discrete Math.* **268** (2003), 343–349.

Received submission date; revised revision date