On classifying objects with specified groups of automorphisms, friendly subgroups, and Sylow tower groups

Leonard H. Soicher

School of Mathematical Sciences

Queen Mary University of London

Mile End Road, London E1 4NS, UK

email: L.H.Soicher@qmul.ac.uk

Abstract

We describe some group theory which is useful in the classification of combinatorial objects having given groups of automorphisms. In particular, we show the usefulness of the concept of a friendly subgroup: a subgroup H of a group K is a *friendly* subgroup of K if every subgroup of K isomorphic to H is conjugate in K to H. We explore easy-to-test sufficient conditions for a subgroup H to be a friendly subgroup of a finite group K, and for this, present an algorithm for determining whether a finite group H is a Sylow tower group. We also classify the maximal partial spreads invariant under a group of order 5 in both $\mathrm{PG}(3,7)$ and $\mathrm{PG}(3,8)$.

1   Introduction

The following situation is common when classifying combinatorial objects. We have some appropriate finite G-set Ω, and we want to classify the elements of Ω satisfying a given G-invariant property, up to the action of the group G.

For example, suppose we are classifying t-(v, k, λ) designs up to isomorphism, for some given t, v, k, λ. Then $G = \mathrm{Sym}(v)$ and Ω can be taken to be the set of all size $b := \lambda\binom{v}{t}/\binom{k}{t}$ multisets of k-subsets of $V := \{1, \ldots, v\}$. The action of G on Ω is the natural one on multisets of subsets of V (the natural action on multisets is like the natural action on sets but with multiplicities preserved), and the property we require of an element $\mathcal{B} \in \Omega$ is that every t-subset of V is contained in exactly λ elements of $\mathcal{B}$; that is, we require that $(V, \mathcal{B})$ is a t-(v, k, λ) design.

Often a complete classification is too difficult, and we impose the additional condition that each object we seek is invariant under at least one of some specified (non-conjugate) subgroups of G. This is the situation we consider in this paper (see also [1], [4], [9], [10], [11], [14]). We are especially interested in the avoidance of unnecessary G-equivalence checks of H-invariant objects (which are often done via the determination of canonical G-orbit representatives of the given objects (see [8], [12], [13])), when often easier $N_G(H)$-equivalence checks suffice (where $N_G(H) = \{x \in G : x^{-1}Hx = H\}$).

Our approach is group-theoretical, and we show the usefulness of friendly subgroups: a subgroup H of a group K is a *friendly* subgroup of K if every subgroup of K isomorphic to H is conjugate in K to H. We note that some of the algebraic approaches in [11] to the classification of t-designs invariant under a given group of automorphisms use Sylow subgroups in a similar way to one way we employ friendly subgroups. See also [10].

After giving illustrative examples (the classification of the maximal partial spreads invariant under a group of order 5 in both $\mathrm{PG}(3, 7)$ and $\mathrm{PG}(3, 8)$), we shall explore easy-to-test sufficient conditions for a subgroup H to be a friendly subgroup of a finite group K, and for this, present an algorithm for determining whether a finite group H is a Sylow tower group.

The theory in this paper is implemented in the DESIGN package [15]. This GAP [2] package can construct, classify, partition and study block designs satisfying a very wide range of user-specified properties, including being invariant under a given group H of automorphisms. The designs may be t-designs (with or without repeated blocks), but in general need not have constant block-size nor constant replication-number.

2 Friendly subgroups

The usefulness of friendly subgroups is based on the following two straight-forward results.

Proposition 2.1 *Let G be a group acting on a set Ω, and let $\mathcal{A}, \mathcal{B} \in \Omega$, with H a friendly subgroup of $G_\mathcal{A}$ (the stabiliser in G of $\mathcal{A}$), and H a subgroup of $G_\mathcal{B}$. Then $\mathcal{A}$ and $\mathcal{B}$ are in the same G-orbit if and only if they are in the same $N_G(H)$-orbit.*

Proof. The if-part is trivial. For the converse, suppose $x \in G$ with $\mathcal{A}^x = \mathcal{B}$. Then $G_\mathcal{B} = (G_\mathcal{A})^x = x^{-1} G_\mathcal{A} x$, and so H^x is a friendly subgroup of $G_\mathcal{B}$. Since $H \leq G_\mathcal{B}$, it must be conjugate in $G_\mathcal{B}$ to H^x, and so there is a $y \in G_\mathcal{B}$ with $H^{xy} = H$. We thus have $xy \in N_G(H)$ and $\mathcal{A}^{xy} = \mathcal{B}^y = \mathcal{B}$. ∎

Proposition 2.2 *Let G be a group acting on a set Ω, and let $\mathcal{A}, \mathcal{B} \in \Omega$, with H a friendly subgroup of $G_\mathcal{A}$. Then if $\mathcal{B}$ is in the same G-orbit as $\mathcal{A}$, every subgroup of $G_\mathcal{B}$ that is isomorphic to H is conjugate in G to H.*

Proof. Suppose $x \in G$ with $\mathcal{A}^x = \mathcal{B}$. Then $G_\mathcal{B} = (G_\mathcal{A})^x$, and so H^x is a friendly subgroup of $G_\mathcal{B}$. Thus, if $J \leq G_\mathcal{B}$ with $J \cong H$, then J is conjugate in $G_\mathcal{B}$ to H^x, and so J is conjugate in G to H. ∎

When classifying H-invariant objects up to G-equivalence (that is, up to being in the same G-orbit), for a given $H \leq G$, one often first classifies the objects up to $N_G(H)$-equivalence (see [1], [15]). Proposition 2.1 allows us to avoid many tests to determine G-equivalence when $N_G(H)$-orbit representatives of the H-invariant objects have already been determined. For such an $N_G(H)$-orbit representative $\mathcal{A}$, if H is a friendly subgroup of $G_\mathcal{A}$ then no G-equivalence tests involving $\mathcal{A}$ are required.

When classifying objects such that each is invariant under at least one of the groups in a given set $\{H_1, \ldots, H_m\}$ of pairwise isomorphic, but non-conjugate, subgroups of G (such as a set of conjugacy class representatives of the subgroups of some prime order p), Proposition 2.2 allows us to avoid many tests to determine when an H_i-invariant object is G-equivalent to an H_j-invariant one. For a given H_i-invariant $\mathcal{A}$, if H_i is a friendly subgroup of $G_\mathcal{A}$, then $\mathcal{A}$ cannot be in the same G-orbit as an H_j-invariant object, when $i \neq j$.

We note that, unlike the methods given in [10], [11], we generally need to be able to directly compute the G-stabilisers (automorphism groups) of the objects we classify, but this avoids certain other computations often required by the methods in [10], [11]. Depending on the problem under consideration, determining G-stabilisers can be easier in practice than determining canonical G-orbit representatives (which is what we aim to avoid having to do). For example, in GAP, when G is a permutation group, determining stabilisers of point-sets is faster in practice than Linton's important procedure [12] to find the lexicographically least set in a G-orbit of a given point-set. Although methods more efficient than Linton's are being developed [8], they still appear to require more work than determining a G-stabiliser of the set under consideration. On the other hand, there appears to be little time difference to computing the automorphism group of a graph and also canonically labelling that graph using nauty [13]. Sometimes, however, it is possible to avoid computing G-stabilisers altogether, for we may be able to deduce that a non-trivial subgroup J of our assumed group of automorphisms H must be a friendly subgroup of any possible automorphism group of the objects we seek. (See [11] for some instances of this and also our illustrative examples in the next section.) Then Proposition 2.1 tells us that to classify our H-invariant objects up to isomorphism, it suffices to classify them up to $N_G(J)$-equivalence.

3   Illustrative examples

To illustrate the application of friendly subgroups, we now describe the classification of the maximal partial spreads invariant under a group of order 5 in both $\mathrm{PG}(3,7)$ and $\mathrm{PG}(3,8)$. Some of the explanation below is taken from the author's webpage [17].

First, some definitions. Let V be an $(n+1)$-dimensional vector space over the finite field $\mathbb{F}_q$. The *projective space* $\mathrm{PG}(n,q)$ is the geometry whose elements are the subspaces of V, with two elements being *incident* if one is contained in the other. The *points* and *lines* of $\mathrm{PG}(n,q)$ are respectively the 1- and 2-dimensional subspaces of V. We identify a line with the set of points it contains. A *partial spread* in $\mathrm{PG}(n,q)$ is a set of lines, no pair of which has a point in common, and a partial spread in $\mathrm{PG}(n,q)$ is *maximal* if it is not properly contained in any partial spread of $\mathrm{PG}(n,q)$. Two partial spreads in $\mathrm{PG}(n,q)$ are *equivalent* if there is an element of the projective general linear

4

group $\mathrm{PGL}(n+1, q)$ (acting naturally on sets of projective lines) mapping one of the partial spreads to the other.

To classify the maximal partial spreads in $\mathrm{PG}(n, q)$ invariant under a given subgroup H of $G := \mathrm{PGL}(n+1, q)$, we form the graph Δ whose vertices are the H-orbits of pairwise disjoint lines, with two vertices of Δ joined by an edge precisely when the vertices are distinct and their union is again a set of pairwise disjoint lines. Then the H-invariant partial spreads correspond to the cliques of Δ, and the H-invariant maximal partial spreads are precisely the unions U of the maximal cliques of Δ, such that no line of $\mathrm{PG}(n, q)$ not in U is disjoint from every line in U. Now the normaliser $N := N_G(H)$ of H acts as a group of automorphisms of Δ, and we use the GRAPE package [16] to construct Δ and determine its maximal cliques, up to the action of N. We then determine which of these cliques gives us a maximal partial spread. The question then is, which elements of the resulting set M of N-inequivalent maximal partial spreads are equivalent under the action of G? If H is a friendly subgroup of each G-stabiliser of a partial spread in M, then we know that the elements of M are pairwise inequivalent. This is in fact the case when $n = 3$, $q = 7$ or 8, and H is a subgroup of order 5.

Now a Sylow 5-subgroup of $G := \mathrm{PGL}(4, 7)$ is cyclic of order 25. It follows that there is a single G-conjugacy class of subgroups of order 5. Let H be such a subgroup, and let $N = N_G(H)$. We have $|G| = 4635182361600$ and $|N| = 1600$.

For $(n, q) = (3, 7)$ we construct the graph Δ on H-orbits of pairwise disjoint lines, as described above, and find that Δ has just 410 vertices. We then find the maximal cliques of Δ, up to N-equivalence, and determine those corresponding to maximal partial spreads of $\mathrm{PG}(3, 7)$, obtaining the set M as above of maximal partial spreads. All this work takes under 26 seconds of CPU-time on a Pentium i5 PC running Linux.

Now if K is any subgroup of G containing H, then H must be a friendly subgroup of K, since H is either a Sylow 5-subgroup of K, or the unique subgroup of order 5 in a Sylow 5-subgroup of K. Thus, M is a set of G-equivalence class representatives of the H-invariant maximal partial spreads. Counting the maximal partial spreads of $\mathrm{PG}(3, 7)$ invariant under a group of order 5 we find that, up to equivalence, there are exactly

$$426, \ 1332, \ 419, \ 94, \ 15, \ 15$$

such maximal partial spreads of respective sizes

$$25, \ 30, \ 35, \ 40, \ 45, \ 50$$

(and no others). An unnecessary (but reassuring) check for G-inequivalence of these maximal partial spreads using Steve Linton's program described in [12] (and available in [16] as the function `SmallestImageSet`) took about 3140 CPU-seconds.

Since a Sylow 5-subgroup of $G := \mathrm{PGL}(4, 8)$ is also cyclic (in fact it is cyclic of order 5), we may classify the maximal partial spreads in $\mathrm{PG}(3, 8)$ invariant under a group H of order 5 as we did in $\mathrm{PG}(3, 7)$. We find that, up to $N_G(H)$-equivalence, and hence up to G-equivalence, there are exactly

$$3, \ 7530, \ 80501, \ 33738, \ 3111, \ 237, \ 18$$

such maximal partial spreads of respective sizes

$$25, \ 30, \ 35, \ 40, \ 45, \ 50, \ 65$$

(and no others). The total time taken for this classification was about 2133 CPU-seconds (without any final unnecessary checks for G-inequivalence). We believe it was previously unknown that $\mathrm{PG}(3, 8)$ has a maximal partial spread having as few as 25 lines (see [7]).

4   Verifying friendliness

Given finite groups H and K, with H a subgroup of K, it is often possible to use relatively cheap computational tests to verify that H is a friendly subgroup of K (when it is such a subgroup). This is often the case when H is a *Hall subgroup* of K, that is, when the order of H is coprime to its index in K. In particular, if H is a Sylow tower group (defined below) and a Hall subgroup of K then H is a friendly subgroup of K (see Theorem 4.1). Typically in our applications, H is small and fixed for many overgroups K.

Suppose H is a finite group, and $1 = T_0 < T_1 < \ldots < T_k = H$ is a normal series for H, such that, for $i = 1, \ldots, k$, T_i/T_{i-1} is isomorphic to a Sylow p_i-subgroup of H, for some prime p_i dividing $|H|$. Then H is called a *Sylow tower group* having *complexion* $(p_1, \ldots, p_k)$ (unlike some authors, we do not require that a Sylow tower group has a complexion $(p_1, \ldots, p_k)$ satisfying $p_1 > \cdots > p_k$). In the next section we give an algorithm to determine

whether a given finite group is a Sylow tower group. Note that if a finite group H is a Sylow tower group then H is soluble, but the converse does not hold in general. However, each finite supersoluble group is a Sylow tower group (see [6]).

The following result details many instances of friendly subgroups of finite groups.

Theorem 4.1 *Let K be a finite group and H a subgroup of K. Then H is a friendly subgroup of K if one or more of the following holds:*

1. *$H = K$;*

2. *K is cyclic;*

3. *H is a Hall subgroup of K and H is a Sylow tower group;*

4. *H is a nilpotent Hall subgroup of K (such as a Sylow subgroup), or more generally, H is a friendly subgroup of a nilpotent Hall subgroup of K;*

5. *K is soluble and H is a Hall subgroup of K, or more generally, K is soluble and H is a friendly subgroup of a Hall subgroup of K;*

6. *H is a simple normal subgroup of K and $|H|^2$ does not divide $|K|$.*

Proof.

1. Trivial.

2. If K is cyclic then H is the unique subgroup of K of order $|H|$.

3. Suppose H is a Hall subgroup of K and is a Sylow tower group of complexion $(p_1, \ldots, p_k)$. Now if $J \leq K$ with $J \cong H$ then J is also a Hall subgroup of K and a Sylow tower group of complexion $(p_1, \ldots, p_k)$, so by [6, Theorem A1], J is conjugate to H.

4. Suppose that H is a friendly subgroup of a nilpotent Hall subgroup L of K, and that $J \leq K$ with $J \cong H$. Then, by [18], J is conjugate to a subgroup J^* of L, and since H is a friendly subgroup of L, J^* (being isomorphic to H) is conjugate to H.

5. Suppose K is soluble, that H is a friendly subgroup of a Hall subgroup L of K, and that $J \leq K$ with $J \cong H$. Then, by [5], J is conjugate to a subgroup J^* of L, and since H is a friendly subgroup of L, J^* (being isomorphic to H) is conjugate to H.

6. Suppose H is a simple normal subgroup of K. Let $J \leq K$ with $H \neq J \cong H$. Since $H \cap J$ is a normal subgroup of J, we must have $H \cap J = 1$, and so HJ is a subgroup of K of order $|H|^2$. Thus, if $|H|^2$ does not divide $|K|$, then H is the only subgroup of its isomorphism class in K.

■

It is worth noting that F. Gross, using the Classification of Finite Simple Groups, shows that every odd-order Hall subgroup of a finite group is a friendly subgroup of that group (see [3]), but I prefer not to use this sledge-hammer to crack the odd nut.

5 Determining whether H is a Sylow tower group

We now present an algorithm which determines whether a given finite group H is a Sylow tower group. The algorithm given here is implemented in GAP and works well in practice for permutation groups and PC-groups (groups with polycyclic presentations), making use of the ability to compute a set of representatives of the Sylow subgroups of a given finite soluble group.

Suppose H is a finite group, and suppose $1 = T_0 < T_1 < ... < T_m \leq H$, such that, for $i = 1, \ldots, m$, T_i is a normal subgroup of H and T_i/T_{i-1} is isomorphic to a Sylow p_i-subgroup of H, for some prime p_i dividing $|H|$. Then $(T_0, \ldots, T_m)$ is called a *partial Sylow tower* for H, having *complexion* $(p_1, \ldots, p_m)$, and a partial Sylow tower for H is a *Sylow tower* for H if its last element is equal to H.

Lemma 5.1 *Suppose $(T_0, \ldots, T_m)$ is a partial Sylow tower, having complexion $(p_1, \ldots, p_m)$, for a finite group H, and let $i \in \{1, \ldots, m\}$. Then $T_i = T_{i-1}P$, for every Sylow p_i-subgroup P of H.*

Proof. Since T_i/T_{i-1} is isomorphic to a Sylow p_i-subgroup of H, the order of T_i is such that T_i contains a Sylow p_i-subgroup of H, and since T_i is normal

8

in H, T_i contains every Sylow p_i-subgroup of H. In addition, since p_i does not divide $|T_{i-1}|$, we conclude that $T_i = T_{i-1}P$, for every Sylow p_i-subgroup P of H. ∎

The following theorem forms the theoretical basis of our iterative algorithm, which, when given a soluble non-nilpotent finite group H, either ascends a Sylow tower for H, or when failing to do so, provides a proof that none such exists. The algorithm is detailed in Figure 1, using the LaTeX package `clrscode`. It would be easy to modify this algorithm to compute a Sylow tower for H when H is a Sylow tower group, since in that case (after removing the nilpotency check and when $H \neq 1$) the i-th iteration of the while-loop of the algorithm constructs in T the i-th term of a Sylow tower, except for the final term, which is H itself.

Theorem 5.2 *Let H be a Sylow tower group, let $p_1, \ldots, p_k$ be the distinct primes dividing $|H|$, and let $\{P_1, \ldots, P_k\}$ be a set of representative Sylow subgroups of H, with P_i a Sylow p_i-subgroup.*

Suppose that $(T_0, \ldots, T_m)$ is a partial Sylow tower for H, with $m < k$. Then there exists an $i \in \{1, \ldots, k\}$ such that T_mP_i is a normal subgroup of H, and properly contains T_m (in which case $(T_0, \ldots, T_m, T_mP_i)$ is a partial Sylow tower for H, extending $(T_0, \ldots, T_m)$).

Proof.   Let $(U_0, \ldots, U_k)$ be a Sylow tower for H, having complexion $(q_1, \ldots, q_k)$ (thus $q_1, \ldots, q_k$ is some particular ordering of $p_1, \ldots, p_k$). Now let j be the least element of $\{1, \ldots, m+1\}$ such that q_j does not divide $|T_m|$ (such a j exists), and define i by $p_i = q_j$.

By the preceding lemma, U_{j-1} is generated by the union of all the Sylow subgroups for all the primes in $\{q_1, \ldots, q_{j-1}\}$, and by the lemma and the definition of j, we see that U_{j-1} is a subgroup of T_m. By the definition of i, $U_j = U_{j-1}P_i$, and so $T_mP_i = T_mU_j$, the product of two normal subgroups of H. Thus T_mP_i is normal in H, and since p_i does not divide $|T_m|$, we have that T_mP_i/T_m is isomorphic to P_i. ∎

5.1   Some counts of Sylow tower groups

The algorithm given in Figure 1 has been implemented as a GAP function, and we use this to determine the number of Sylow tower groups amongst groups

Figure 1: Algorithm to determine whether H is a Sylow tower group

IsSylowTowerGroup(H)

    ▷ Given a finite group H, this function returns TRUE
    ▷ if H is a Sylow tower group, and returns FALSE if not.
    if H is not soluble
        then return FALSE;
    if H is nilpotent
        then return TRUE;
    $P \leftarrow$ a list of representatives of the non-trivial Sylow subgroups of H;
    ▷ P contains exactly one Sylow p-subgroup for each prime p dividing $|H|$.
    $T \leftarrow \{1_H\}$;
    while Length(P) > 1
            do ▷ T is the highest term of a partial Sylow tower for H,
                ▷ $T \neq H$, and $|T|$ is coprime to the order of each element of P.
                if there is no Q in P such that $\langle T, Q \rangle$ is normal in H
                    then ▷ By Theorem 5.2, H cannot be a Sylow tower group.
                        return FALSE;
                $Q \leftarrow$ the first Q in P such that $\langle T, Q \rangle$ is normal in H;
                $T \leftarrow \langle T, Q \rangle$;
                Remove Q from P;
    ▷ There is now just one element left in P, T is the last term
    ▷ of a partial Sylow tower for H, and this partial Sylow tower
    ▷ can be extended to a Sylow tower by adjoining H itself.
    return TRUE;

of small order, transitive permutation groups of low degree, and primitive permutation groups of modest degree. All timings are in CPU-seconds on a Pentium i5 PC running Linux.

Using the GAP library of small groups, we obtain a list of the 7012 groups (up to isomorphism) of order at most 255. We determine that 6836 of these are Sylow tower groups in about 4 seconds. This compares with 6998 of the groups being soluble, 6590 being supersoluble, and 3722 being nilpotent.

Using the GAP library of transitive groups, we obtain a list of the 4952 transitive groups (up to permutation isomorphism) of degrees 2 to 23. We determine that 2772 of these are Sylow tower groups in about 95 seconds. This compares with 4212 of the groups being soluble, 1819 being supersoluble, and 1487 being nilpotent.

Using the GAP library of primitive groups, we obtain a list of the 2269 primitive groups (up to permutation isomorphism) of degrees 2 to 256. We determine that 702 of these are Sylow tower groups in about 22 seconds. This compares with 939 of the groups being soluble, 438 being supersoluble, and 54 being nilpotent.

6   Acknowledgements

References

[1] Z. Eslami, Classification of designs with nontrivial automorphism groups, *J. Combin. Des.* **14** (2006), 479–489.

[2] The GAP Group, GAP — Groups, Algorithms, and Programming, Version 4.8.3, 2016, `http://www.gap-system.org/`

[3] F. Gross, Conjugacy of odd order Hall subgroups, *Bull. Lond. Math. Soc.* **19** (1987), 311–319.

[4] E. Haberberger, A. Betten and R. Laue, Isomorphism classification of t-designs with group theoretical localisation techniques applied to some

Steiner quadruple systems on 20 points, *Congr. Numer.* **142** (2000), 75–96.

[5] P. Hall, A note on soluble groups, *J. Lond. Math. Soc.* **3** (1928), 98–105.

[6] P. Hall, Theorems like Sylow's, *Proc. Lond. Math. Soc. (3)* **6** (1956), 286–304.

[7] M. Iurlo and S. Rajola, A computer search of maximal partial spreads in $PG(3, q)$, arXiv:1011.5338 (2011), 10 pp.

[8] C. Jefferson, E. Jonauskyte, M. Pfeiffer and R. Waldecker, Minimal and canonical images, arXiv:1703.00197 (2017), 32 pp.

[9] P. Kaski, Isomorph-free exhaustive generation of designs with prescribed groups of automorphisms, *SIAM J. Discrete Math.* **19** (2005), 664–690.

[10] R. Laue, Constructing objects up to isomorphism, simple 9-designs with small parameters, in *Algebraic Combinatorics and Applications* (A. Betten et al., eds), Springer, Berlin-Heidelberg, 2001, 232–260.

[11] R. Laue, Solving isomorphism problems for t-designs, in *Designs 2002: Further computational and constructive design theory* (W.D. Wallis, ed.), Kluwer, Boston, 2003 (Paperback version: Springer, New York, 2012), 277–300.

[12] S. Linton, Finding the smallest image of a set, in *ISSAC 2004: Proceedings of the 2004 International Symposium on Symbolic and Algebraic Computation* (J. Gutierrez, ed.), ACM Press, New York, 2004, 229–234.

[13] B.D. McKay and A. Piperno, Practical graph isomorphism, II, *J. Symbolic Comput.* (2013) **60**, 94–112.

[14] L.H. Soicher, On the structure and classification of SOMAs: generalizations of mutually orthogonal Latin squares, *Electron. J. Combin.* **6** (1999), #R32, 15 pp.

[15] L.H. Soicher, The DESIGN package for GAP, Version 1.6, 2011, `http://designtheory.org/software/gap_design/`

[16] L.H. Soicher, The GRAPE package for GAP, Version 4.7, 2016, `http://www.maths.qmul.ac.uk/~leonard/grape/`

[17] L.H. Soicher, Computation of Partial Spreads,
`http://www.maths.qmul.ac.uk/~leonard/partialspreads/`

[18] H. Wielandt, Zum Satz von Sylow, *Math. Z.* **60** (1954), 407–408.