London Mathematical Society ISSN 1461–1570

ALGORITHMIC RECOGNITION OF ACTIONS OF 2-HOMOGENEOUS GROUPS ON PAIRS

GRAHAM R. SHARP

Abstract

We give an algorithm that takes as input a transitive permutation group (G, Ω) of degree $n = \binom{m}{2}$, and decides whether or not Ω is G-isomorphic to the action of G on the set of unordered pairs of some set Γ on which G acts 2-homogeneously. The algorithm is constructive: if a suitable action exists, then one such will be found, together with a suitable isomorphism. We give a deterministic $O(sn\log^c n)$ implemention of the algorithm that assumes advance knowledge of the suborbits of (G, Ω). This leads to deterministic $O(sn^2)$ and Monte-Carlo $O(sn\log^c n)$ implementations that do not make this assumption.

1. *Introduction*

Let (G, Γ) be a finite permutation group, and define

$$\Gamma^{\{2\}} = \{\{\alpha, \beta\} \mid \alpha, \beta \in \Gamma,\ \alpha \neq \beta\}$$

and

$$\Gamma^{(2)} = \{(\alpha, \beta) \mid \alpha, \beta \in \Gamma,\ \alpha \neq \beta\}.$$

If G is transitive on $\Gamma^{(2)}$, it is described as 2-*transitive* on Γ; if it is transitive only on $\Gamma^{\{2\}}$, then it is described as 2-*homogeneous* on Γ.

Recall that two G-sets Γ and Δ are said to be G-isomorphic if and only if there is a bijection $\eta : \Gamma \to \Delta$ such that $\gamma\eta^g = \gamma^g\eta$ for all $\gamma \in \Gamma$ and all $g \in G$.

The purpose of this paper is to describe an algorithm that takes as input a transitive permutation group (G, Ω) and determines whether or not there exists an action of G on a set Γ such that Ω and $\Gamma^{\{2\}}$ are G-isomorphic. We will call such a set Γ a *solution* to the *exterior square-root problem* for (G, Ω). More precisely, a solution will be a pair (Γ, η) where Γ is a G-set and η a G-isomorphism from $\Gamma^{\{2\}}$ to Ω. The problem is specified as follows.

Specification 1.1.

Input *A transitive permutation group* (G, Ω) *of degree* $n = \#\Omega > 1$, *where* $n = \binom{m}{2}$ *for some integer* m, *given by generators* $g_1, \ldots, g_s \in \mathrm{Sym}(\Omega)$.

Output *A value* $b \in \{\textsc{True}, \textsc{False}\}$.

If $b = \textsc{True}$, *then a* G*-set* Γ *of size* m, *and a bijection* $\eta : \Gamma^{\{2\}} \to \Omega$ *are also returned, satisfying* $\{\gamma_1^g, \gamma_2^g\}\eta = (\{\gamma_1, \gamma_2\}\eta)^g$ *for all* $g \in G$, $\{\gamma_1, \gamma_2\} \in \Gamma^{\{2\}}$.

If $b = \textsc{False}$ *then there does not exist a pair* (Γ, η) *satisfying these conditions.*

The author was supported by the EPSRC.
Received 27th August 1997, revised 17th June 1998*; published 9th November 1998.*
1991 Mathematics Subject Classification 20B40,68Q40

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

The main result of this paper is an algorithm that satisfies Specification 1.1, and that can be implemented in $O(sn^2)$ time.

Note that for the purposes of complexity analysis, we take our standard operation to be finding the image of a given point in Ω under the action of a given permutation in $\mathrm{Sym}(n)$ (as is usual in computational permutation group theory), even though this involves manipulations with integers that require $\log n$ bits to store.

We first use some group theory to classify the permutation groups (G, Γ) that can give rise to solutions; *i.e.*, those groups that are transitive on $\Gamma^{\{2\}}$. The structure of the remainder of the paper is based on the different classes of groups arising out of this taxonomy.

2. *Classifying* 2*-homogeneous groups*

Lemma 2.1. *Suppose (G, Γ) is a permutation group, and that G is transitive on $\Gamma^{\{2\}}$ and that $\#G$ is divisible by* 2*. Then G is transitive on $\Gamma^{(2)}$.*

Proof. G contains an involution t, which must interchange two points α, β of Γ. Let γ, δ be arbitrary points of Γ. Since G is transitive on $\Gamma^{\{2\}}$, there is $g \in G$ with $\{\alpha, \beta\}^g = \{\gamma, \delta\}$. Now either g or tg maps (α, β) to (γ, δ) as ordered pairs, and so G is 2-transitive on Γ. □

Theorem 2.2. *Suppose (G, Γ) is a permutation group, and that G is transitive on $\Gamma^{(2)}$. Then one of two cases arises:*

(i) *the group G has a normal subgroup S which is non-abelian simple, and $C_G(S) = 1$, so the action of G on S by conjugation gives an embedding of G into $\mathrm{Aut}(S)$, or*

(ii) *the group G contains an elementary abelian normal subgroup V which is regular on Γ.*

In each case, the subgroup (S or V) is the unique minimal normal subgroup of G.

If G is transitive on $\Gamma^{\{2\}}$ but not on $\Gamma^{(2)}$ then G satisfies the conditions for case (*ii*) *above.*

Proof. For 2-transitive groups the theorem is a result of Burnside [2, §154, Theorem XIII].

If (G, Γ) is 2-homogeneous but not 2-transitive, then G is of odd order, by the lemma above. By the Odd Order Theorem of Feit and Thompson [5], G is solvable. Also, G is clearly primitive, and a minimal normal subgroup of a solvable primitive group is both elementary abelian and regular [9, Theorem 11.5]. □

In Theorem 2.2, if the minimal normal subgroup of G is simple non-abelian, we say that G is *almost simple*, while if the minimal normal subgroup of G is elementary abelian, we say that G is of *affine type*. If G has an elementary abelian regular normal subgroup V, then its degree is a prime power, p^d say, and V can be regarded as a vector space over $\mathbb{F}_p$. We will write V additively.

Proposition 2.3. *Suppose a permutation group (G, Γ) has an elementary abelian regular normal subgroup V of order p^d. Then G embeds in the group*

$$\mathrm{AGL}(d, p) = \mathrm{AGL}(V) = \{x \mapsto x^h + v \mid h \in \mathrm{GL}(V), v \in V\}$$

of affine transformations of V, with V mapping onto the translation subgroup $T = \{t_v \mid v \in V\}$, where t_v is the map which sends x to $x + v$. The permutation group (G, V) with the action of G on V arising from this embedding of G in $\mathrm{AGL}(V)$ is isomorphic to the group (G, Γ).

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

Proof. See, for example, [6, §2]. □

The structure of the algorithm in this paper is a series of subroutines, each solving a problem of the following form.

Specification 2.4. *Let $\mathcal{C}$ be a class of 2-homogeneous permutation groups.*

Input *A transitive permutation group (G, Ω) of degree $n = \#\Omega > 1$, where $n = \binom{m}{2}$ for some integer m, given by generators $g_1, \ldots, g_s \in \mathrm{Sym}(\Omega)$.*

Output *A value $b \in \{\textsc{True}, \textsc{False}\}$.*

If $b = \textsc{True}$, then a G-set Γ of size m, and a bijection $\eta : \Gamma^{\{2\}} \to \Omega$ are also returned, satisfying $\{\gamma_1^g, \gamma_2^g\}\eta = (\{\gamma_1, \gamma_2\}\eta)^g$ for all $g \in G$, $\{\gamma_1, \gamma_2\} \in \Gamma^{\{2\}}$.

If $b = \textsc{False}$ then there does not exist a pair (Γ, η) that satisfies these conditions, and where (G, Γ) lies in the class $\mathcal{C}$.

Notice that if $\mathcal{C}$ is the whole class of 2-homogeneous permutation groups, then this is exactly equivalent to Specification 1.1. If we have a finite number of subroutines $\mathsf{A}_1, \ldots, \mathsf{A}_k$ satisfying Specification 2.4 for different classes $\mathcal{C}_1, \ldots, \mathcal{C}_k$ respectively such that the union of the $\mathcal{C}_i$ contains the class of all 2-homogeneous permutation groups, then the following procedure satisfies Specification 1.1.

```
1. For $i := 1$ to $k$ do
2.     Call $\mathsf{A}_i$ with input $(G, \Omega)$ and output $b, \Gamma, \eta$.
3.     If $b = \textsc{True}$ then Exit with output $b, \Gamma, \eta$.
4. End for.
5. Exit (with output $b$, which is $\textsc{False}$).
```

In practice, we will use Theorem 2.2 and the classification of finite 2-transitive groups to obtain our classes $\mathcal{C}_i$. The primary division will be into four classes, one $\mathcal{A}_o$ containing all 2-homogeneous affine groups over fields of odd characteristic (by Theorem 2.2 this contains all the 2-homogeneous, non-2-transitive groups); another $\mathcal{A}_e$ containing the 2-transitive affine groups over fields whose size is even and strictly larger than 2; a third $\mathcal{L}$ containing the almost simple groups of Lie type whose socle is $\mathrm{PSU}(3, q)$, $\mathrm{Sz}(q)$ or $\mathrm{R}(q)$ for some q; and a fourth class $\mathcal{Z}$ containing all the other 2-transitive groups. The class $\mathcal{Z}$ will be subdivided further according to the families of groups in the classification of 2-transitive groups. The formal definitions of the main classes and the subclasses of $\mathcal{Z}$ will be given in Table 1 on page 120.

None of the results in this paper are dependent on the classification of finite 2-transitive groups (or the classification of finite simple groups), except the result which says that the union of the four classes $\mathcal{A}_o$, $\mathcal{A}_e$, $\mathcal{L}$ and $\mathcal{Z}$ contains all 2-homogeneous permutation groups.

The organisation of the rest of this paper is loosely based around this subdivision into classes of groups. In Sections 3 to 6, we will give a largely combinatorial solution for the class $\mathcal{Z}$, and in the process of doing this we will develop some combinatorial techniques that will be useful later. In Sections 7 and 8 we will cover the affine classes, $\mathcal{A}_o$ and $\mathcal{A}_e$, and in Section 9 we will cover the one remaining class, $\mathcal{L}$. The final section describes results obtained from a GAP implementation of these techniques.

3. *Solution subsets*

In this section we characterize all solutions as arising in a particular way from a subset of Ω (a so-called 'solution subset') satisfying certain properties. We give an algorithm to check whether a given subset of Ω satisfies these properties, and, if so, to calculate a solution (Γ, η) from it.

Definition 3.1. A *solution subset* of Ω is defined to be a subset A of Ω, of size $m-1$, such that (i) $\#A^G = m$ (where A^G is the orbit containing A in the natural action of G on the power set of Ω) and (ii) for all $g \in G$, either $A = A^g$ or $\#(A \cap A^g) = 1$. In this situation, define the map $\mu : (A^G)^{\{2\}} \to \Omega$ by $\{X, Y\}\mu = \omega$ where $X \cap Y = \{\omega\}$.

Proposition 3.2. *Let A be a solution subset of Ω. Then (A^G, μ) is a solution pair (as defined at the very beginning of the paper), and for any solution (Γ, η), there exists a solution subset A and corresponding map μ, and a G-isomorphism $\tau : \Gamma \to A^G$ such that $\eta = \tau\mu$ (where τ is lifted to map $\Gamma^{\{2\}}$ to $(A^G)^{\{2\}}$).*

Proof. We show first that (A^G, μ) is a solution. To show this, we need to show that μ is a bijection, and that it preserves the action of G. There exists g_0 such that $A^{g_0} \neq A$, and so $A^{g_0} \cap A = \{\omega_0\}$, say. Let $\omega \in \Omega$. Since G is transitive on Ω, there is $g \in G$ with $\omega_0^g = \omega$, and now $\{A^g, A^{g_0 g}\}\mu = \omega$, so μ is surjective. Since $(A^G)^{\{2\}}$ and Ω have the same size, this means that μ is a bijection. Let $g_1, g_2, g \in G$. Then

$$\{\{A^{g_1 g}, A^{g_2 g}\}\mu\} = A^{g_1 g} \cap A^{g_2 g} = (A^{g_1} \cap A^{g_2})^g = \{(A^{g_1} \cap A^{g_2})\mu\}^g$$

and so (A^G, μ) is a solution.

Now suppose that (Γ, η) is a solution. For $\gamma \in \Gamma$, define

$$A_\gamma = \left\{\{\gamma, \gamma'\}\eta \;\middle|\; \gamma' \in \Gamma \setminus \{\gamma\}\right\},$$

and for some $\gamma_0 \in \Gamma$, take $A = A_{\gamma_0}$. Certainly $A_\gamma^g = A_{\gamma^g}$ for all $\gamma \in \Gamma$ and all $g \in G$, since η is a G-isomorphism, so $A^G = \{A_\gamma \mid \gamma \in \Gamma\}$. Therefore A^G has size m; it is clear from the definition that A has size $m-1$. If $\gamma_1, \gamma_2 \in \Gamma$ are distinct then we have

$$A_{\gamma_1} \cap A_{\gamma_2} = \{\{\gamma_1, \gamma_2\}\eta\} \tag{1}$$

which shows that the intersection of distinct elements of A^G has size 1, as required, so A is a solution subset. Define τ by $\gamma\tau = A_\gamma$; it is clear that this is a G-isomorphism. It is now evident from (1) and the definition of μ that $\eta = \tau\mu$. □

This means that we may restrict our search for solutions to the exterior square-root problem for (G, Ω) to a search for solution subsets.

Lemma 3.3. *Let A be a solution subset. Then for each $\omega \in \Omega$ there are precisely two sets $X \in A^G$ such that $\omega \in X$.*

Proof. Certainly there at least two sets X containing ω, since the map μ is surjective. However the number of elements in the union of the sets in A^G *counting repetitions* is only $m(m-1) = 2\#\Omega$. Therefore there are exactly two such sets containing ω. □

We now give an algorithm for deciding whether a given subset $A \subseteq \Omega$ is a solution subset.

Proposition 3.4. *There is an algorithm which given $A \subset \Omega$ of size $m-1$ will decide whether A is a solution subset, and find a solution pair (Γ, η) if it is, in $O(sn)$ time.*

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

Proof. The algorithm is shown as Algorithm 3.5. It takes (G, Ω) and the putative solution subset A as input, and outputs b, Γ, η as in Specification 1.1, except that here if $b = \textsc{False}$ then we only know that A is not a solution subset, rather than that no solution exists.

The structure of the algorithm is basically that of a traversal of the graph with the points of A^G as vertices, and an edge (X, Y) for each of the original generators of G that maps X to Y. However, there are certain extra features that make use of the extra structure of the problem.

Algorithm 3.5. TestSolutionSubset(A)

```
1. Initialize $W(\omega) := \emptyset$ for all $\omega \in \Omega$.
2. For each $\omega \in A$, set $W(\omega) := \{A\}$.
3. Initialize $Q := \{A\}$, $T := \{A\}$, $b := \textsc{False}$.
```

Here T stores the sets in A^G as they are found, and Q contains those sets in T that have been found, but whose neighbours have not yet been checked. For each $\omega \in \Omega$, the set $W(\omega)$ will contain those elements of T (that is, of that part of the orbit A^G so far discovered) that contain ω.

At the start of the main loop, we choose a point X that has not been fully processed, and process it by considering its neighbours:

```
4. While $Q \neq \emptyset$ do
5.     Choose $X \in Q$, and set $Q := Q \setminus \{X\}$.
6.     For $g$ in $\{g_1, \ldots, g_s\}$ do
7.         Form $S := X^g$.
8.         Choose $\omega_1, \omega_2 \in S$.
9.         If $W(\omega_1) = \emptyset$ or $W(\omega_2) = \emptyset$ or $(\#W(\omega_1) = \#W(\omega_2) = 1$ and $W(\omega_1) \neq$
               $W(\omega_2))$ then
```

If the test in line 9 is passed then S is a new element of A^G. If it fails, then either S is not a new element of A^G, or it is new, but there is another element of A^G, which has already been discovered, whose intersection with S is not of size 1.

If the test is passed, we update Q and T; if we have generated more than m elements of A^G then we know A is not a solution subset. We then update W:

```
10.            Set $Q := Q \cup \{S\}$ and $T := T \cup \{S\}$; if $\#T > m$ then Exit.
11.            Set $U := \emptyset$.
12.            For $\omega \in S$ do
13.                If $\#W(\omega) > 1$ or $W(\omega) \cap U \neq \emptyset$ then Exit
14.                Else
15.                    Set $U := U \cup W(\omega)$.
16.                    Set $W(\omega) := W(\omega) \cup \{S\}$.
17.                End if.
18.            End do.
```

We know that no $W(\omega)$ should ever contain more than two elements, and that no two $W(\omega)$ that both contain two elements should ever be the same. If either of these situations arises we can immediately deduce that A is not a solution subset. If neither of these situations arises in any of the $W(\omega)$ corresponding to those $\omega \in S$, then we can update W by adding S to these $W(\omega)$. This concludes the processing in the case where we have a possible new member of S, and we now turn to the other case:

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

```
19.             Else if $\# \bigcap_{\omega \in S} W(\omega) \neq 1$ then Exit.
20.             End if.
```

This ‘else’ matches the ‘if’ in line 9, so if this code is executed then we know that either S is already in T, or it is not but either has intersection of size greater than 1 with an element of T, or contains a point that already lies in two different elements of T. The test in this line determines whether S is already in T; if not, then the program concludes that A is not a solution set; otherwise, we continue, ready to consider a new element S.

The loops now repeat, and so we traverse the orbit A^G in the usual way.

```
21.         End for.
22.     End while.
23.     If $\#T = m$ then
24.         Set $b :=$ TRUE, $\Gamma := T$ and define $\eta : T^{\{2\}} \to \Omega$ by $\{X_1, X_2\}\eta := \omega$, where
                $\{\omega\} = X_1 \cap X_2$.
25.     End if.
26.     Exit.
```

If the loops are completed then $T = A^G$, so if the condition in the final ‘if’ statement is true then A^G has size m, and $m - 1$ copies of the m elements of A^G have been formed into $m(m-1)/2$ sets $W(\omega)$ in such a way that no $W(\omega)$ contains more than two elements, and every $W(\omega)$ is different. Therefore all the $W(\omega)$ are pairs, and

$$\{W(\omega) \bigm| \omega \in \Omega\} = (A^G)^{\{2\}}.$$

For any $X \in T$, $X = \{\omega \in \Omega \bigm| X \in W(\omega)\}$, and therefore $A \cap A^g$ is either A itself or has size 1, since $T = A^G$, and so A is a solution set. It now follows from the earlier results on solution sets that (Γ, η) is a solution, where Γ and η are as defined in the algorithm.

When implementing this algorithm, we can number the elements of T as they are created, and use these index numbers, instead of the elements of T themselves, as the elements of W, U and Q. Under the (standard) assumption that manipulations with natural numbers of magnitude $O(n)$ can be made in constant time (see the beginning of the paper), this enables us to implement Q so that all references to it require only constant time, to do the test in line 9 in constant time, and the test in line 19 in $O(m)$ time (recall that each $W(\omega)$ has size at most 2). Making S can be done in $O(m)$ time, as can adding S to T. This is also the case for the inner ‘For’ loop, if we implement the set U as a look-up table indexed by the numbers $1, \ldots, k$, so determining whether a particular number lies in U takes constant time. The main loop (lines 7–20) is executed no more than sm times. Thus the overall complexity of the main loop is $O(sm^2)$, that is, $O(sn)$.

Extending the idea of the index numbers, we can return $\Gamma = \{1, \ldots, m\}$; if $g \in G$ then the image of $i \in \Gamma$ under the action of g can be found by taking two points ω_1, ω_2 in the set in T corresponding to i, and then i^g is the one point in the intersection of $W(\omega_1^g)$ and $W(\omega_2^g)$. Calculating the images of the generators in this way requires $O(sm)$ time.

Computing the map η still requires only $O(n)$ time, as we can do it in reverse: to create an array indexed by $\Gamma^{\{2\}}$ containing the image of that pair under the mapping η, we can take each $\omega \in \Omega$, read off the corresponding pair of indices from $W(\omega)$, and enter ω in the appropriate place in the array.

Therefore the given algorithm can be implemented in $O(sn)$ time. □

4. *The adjacent-point set*

We now introduce another type of subset of Ω, which also characterizes solutions, and which is easier to find in practice. We give an algorithm for testing whether a given subset is of this type and, if so, for constructing a solution.

Definition 4.1. Suppose (Γ, η) is a solution, and $\omega \in \Omega$. Define

$$\Lambda_\Gamma(\omega) = \left\{\omega' \in \Omega \;\middle|\; \omega'\eta^{-1} \cap \omega\eta^{-1} \neq \emptyset\right\}.$$

Thus $\Lambda_\Gamma(\omega)$ is the set of the images under η of all those pairs in $\Gamma^{\{2\}}$ that are 'adjacent' to (have non-trivial intersection with) the pre-image of ω. This is a G_ω-invariant set of size $2m - 3$, containing ω. The subscript Γ is used to indicate the dependence of Λ_Γ on the solution (Γ, η); perhaps better, but cumbersome, would be to include η in the notation as well. When a solution is given in terms of a solution subset A, so that $\Gamma = A^G$ and $\eta = \mu$, we will use Λ_A instead of Λ_{A^G}. If A is a solution set, then there are group elements g_1 and g_2 such that $\{\omega\} = A^{g_1} \cap A^{g_2}$, and then it is easily checked that $\Lambda_A(\omega) = A^{g_1} \cup A^{g_2}$. Clearly, $\Lambda_\Gamma(\omega)^g = \Lambda_\Gamma(\omega^g)$ for all $g \in G$.

Now suppose we have a G_ω-invariant subset L of Ω which contains ω and has size $2m - 3$. We can decide whether or not there is a solution (Γ, η) such that $\Lambda_\Gamma(\omega) = L$, and find one if one exists, as follows.

Algorithm 4.2. TestSet(ω, L)

Input *Generators for a transitive permutation group* (G, Ω) *of degree* $\#\Omega = n = \binom{m}{2}$
where $m > 4$,
a point $\omega \in \Omega$, *and*
a G_ω*-invariant subset* L *of* Ω *of size* $2m - 3$ *that contains* ω.

Output *A value* $b \in \{\text{True}, \text{False}\}$.
If $b = \text{True}$, *then a* G*-set* Γ *of size* m, *and a bijection* $\eta : \Gamma^{\{2\}} \to \Omega$ *are also returned, satisfying* $\{\gamma_1^g, \gamma_2^g\}\eta = (\{\gamma_1, \gamma_2\}\eta)^g$ *for all* $g \in G$, $\{\gamma_1, \gamma_2\} \in \Gamma^{\{2\}}$, *and with* $\Lambda_\Gamma(\omega) = L$.
If $b = \text{False}$ *then there does not exist a pair* (Γ, η) *that satisfies these conditions.*

1. Construct a set $\mathcal{T}_1$ of translates of L, stopping when two elements L_1 and L_2 of $\mathcal{T}_1$ intersect in a set of size precisely m, or when $\#\mathcal{T}_1 > m/2$; if this limit is exceeded without finding L_1 and L_2 then Return $b := \text{False}$.
2. Set $P := L_1 \cap L_2$.
3. Construct a set $\mathcal{T}_2$ of translates of P, stopping when two elements P_1 and P_2 of $\mathcal{T}_2$ intersect in a set of size precisely $m - 1$, or when $\#\mathcal{T}_2 > m$; if this limit is exceeded without finding P_1 and P_2 then return $b := \text{False}$.
4. Set $B := P_1 \cap P_2$.
5. TestSolutionSubset(B).

The algorithm requires $O(sn)$ time.

Proof. We prove that if there exists a solution subset A such that $\Lambda_A(\omega) = L$, then the set B computed by the algorithm is a solution subset, and $\Lambda_B(\omega) = L$. It then follows that the call to TestSolutionSubset yields the correct answer. So suppose that such a subset A exists. We can take Γ to be the set A^G of translates of A. We have $\Lambda_A(\omega^g) = L^g$ for all $g \in G$, so each translate of L is an adjacent-point set.

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

We consider the intersection of distinct adjacent-point sets. Suppose $\Lambda_A(\omega_i) = A^{g_i} \cup A^{g_i'}$ for $i = 1, 2$. Then

$$\begin{aligned}\Lambda_A(\omega_1) \cap \Lambda_A(\omega_2) &= (A^{g_1} \cup A^{g_1'}) \cap (A^{g_2} \cup A^{g_2'}) \\ &= (A^{g_1} \cap A^{g_2}) \cup (A^{g_1} \cap A^{g_2'}) \cup (A^{g_1'} \cap A^{g_2}) \cup (A^{g_1'} \cap A^{g_2'}).\end{aligned}$$

As the $\Lambda_A(\omega')$ are distinct and A^{g_i} is distinct from $A^{g_i'}$ for $i = 1, 2$, only two cases arise: either all the four sets A^{g_i}, $A^{g_i'}$ are different, or there is one pair the same and the others are different; in that case we may without loss of generality assume that $A^{g_1} = A^{g_2}$.

If all four sets are different then by Lemma 3.3, the intersection $\Lambda_A(\omega_1) \cap \Lambda_A(\omega_2)$ has size 4. If $A^{g_1} = A^{g_2}$ then

$$\Lambda_A(\omega_1) \cap \Lambda_A(\omega_2) = A^{g_1} \cup (A^{g_1'} \cap A^{g_2'}),$$

which has size m. Since $m > 4$ the algorithm can distinguish between these two cases, and the intersection of L_1 and L_2 falls in the second case. Note that there are m translates of A in Γ, and each translate of L contains two translates of A. If the pairwise intersections of all the elements of $\mathcal{T}_1$ fall in the first case above, then the translates of A involved in the elements of $\mathcal{T}_1$ must all be different, so in this case $\mathcal{T}_1$ can have size at most $m/2$. Therefore, if a pair L_1, L_2 exists, we must have found such a pair by the time we are processing the $(\lfloor m/2 \rfloor + 1)$-th distinct translate of L.

For the next step, we can assume that $P = A^{h_1} \cup (A^{h_2} \cap A^{h_3})$ for some distinct A^{h_1}, A^{h_2}, A^{h_3}. Translates of P are also of this form. Let $A^{h_1} \cup (A^{h_2} \cap A^{h_3})$ and $A^{h_1'} \cup (A^{h_2'} \cap A^{h_3'})$ be distinct translates of P, and consider their intersection. If $A^{h_1} = A^{h_1'}$ then the intersection contains A^{h_1}; as the two translates are distinct and the size of A^{h_1} is only 1 less than that of P, this must be the whole of the intersection. Thus in this case the intersection has size $m - 1$ and is a solution set, as it is a translate of A; also $\Lambda_{A^{h_1}}(\omega) = \Lambda_A(\omega) = L$.

If $A^{h_1} \neq A^{h_1'}$ then their intersection has size 1; the only other points that could possibly lie in the intersection of the two translates of P are the points in the two singleton sets $A^{h_2} \cap A^{h_3}$ and $A^{h_2'} \cap A^{h_3'}$, so the intersection of these two translates of P has size at most 3, and as $m > 4$ the algorithm can recognise when the intersection is of size $m - 1$. Observe that if the set $\mathcal{T}_2$ of translates of P is such that none of the pairwise intersections of elements of $\mathcal{T}_2$ has size $m - 1$, then the size of $\mathcal{T}_2$ is at most m, since each element of $\mathcal{T}_2$ must contain a different translate of A.

To implement the algorithm in the stated time bound we use a method similar to that of Algorithm 3.5. Consider the calculation of $\mathcal{T}_1$, which is shown in more detail in Algorithm 4.3 on page 117. For each point ξ of Ω we store a list $W(\xi)$ of elements of $\mathcal{T}_1$ containing that point. To process a new translate S of L we use a list C indexed by $\mathcal{T}_1$, whose entries are initially 0. For each point ξ of S, we find the list of elements of $\mathcal{T}_1$ containing ξ, and for each element T of that list, increment $C(T)$. We thus find out the size of the intersection of S with each element of $\mathcal{T}_1$. We use index numbers to represent the elements of $\mathcal{T}_1$, as in the implementation of Algorithm 3.5. This enables W, Q and C to be manipulated quickly. The outer two loops execute $O(ms)$ times, as $\#\mathcal{T}_1$ is $O(m)$. We show that the body of the inner of these two loops, from line 7 to line 25, always executes in $O(m)$ time, except during the last iteration, when it executes in $O(m^2)$ time. This will be clear if we can show that $\sum_{T \in \mathcal{T}} \#(S \cap T)$ is $O(m)$ or $O(m^2)$ respectively, since all lines apart from the nested loops calculating C clearly run in $O(m)$ time (if S is added to $\mathcal{T}$ in line 15, then it will be assigned an index number that will enable the subsequent updating of $W(\xi)$ to be completed in constant time for each $\xi \in S$).

Algorithm 4.3. Implementation of Step 1 of Algorithm 4.2

```
Initialize $W(\xi) := \emptyset$ for all $\xi \in \Omega$;
For each $\xi \in L$ set $W(\xi) := \{L\}$;
Initialize $Q := \{L\}$, $\mathcal{T}_1 := \{L\}$;
While $Q \neq \emptyset$ do
    Choose $X \in Q$, and set $Q := Q \setminus \{X\}$;
    For $g$ in $\{g_1, \ldots, g_s\}$ do
        Form $S := X^g$;
        Initialize $C$ to be the zero function from $\mathcal{T}_1$ to $\mathbb{N}$;
        For each $\xi$ in $S$ do
            For each $T$ in $W(\xi)$ do
                Increment $C(T)$;
            End for;
        End for;
        If $C(T) = 4$ for all $T \in \mathcal{T}_1$ then
            Add $S$ to $Q$ and to $\mathcal{T}_1$;
            If $\#\mathcal{T}_1 > m/2$ then exit with $b :=$ FALSE;
            For $\xi$ in $S$ do
                Add $S$ to $W(\xi)$;
            End for;
        Else if there exists $T \in \mathcal{T}_1$ such that $C(T) = m$ then
            Exit with $L_1, L_2 := S, T$;
        Else if there exists $T \in \mathcal{T}_1$ such that $C(T)$ is not 4, $m$ or $2m - 3$ then
            Exit with $b :=$ FALSE;
        Else do nothing (* $S$ is already in $\mathcal{T}_1$ *);
        End if;
    End for;
End while;
End (* This point will not be reached—one of the exit conditions above will be
        satisfied first *).
```

Certainly $\sum_{T \in \mathcal{T}} \#(S \cap T)$ is $O(m^2)$ since each intersection has size at most $2m - 3$. If this iteration of the main loop body is not the last then: every intersection has size 4, m or $2m - 3$; at most one has size $2m - 3$ (as all entries in $\mathcal{T}_1$ are distinct); and none has size m. Thus $\sum_{T \in \mathcal{T}} \#(S \cap T)$ is at most $2m - 3 + 4\#\mathcal{T}_1$, and so is $O(m)$.

A similar method is used to find P_1 and P_2 in the same time bound, and the other steps of the algorithm can be implemented in $O(sm^2)$ time. The result follows, since $O(sm^2)$ is the same as $O(sn)$. □

Observe that we have not used the fact that L is G_ω-invariant at any point. However, we will later use this requirement to restrict the possible candidates for the set L.

The following variant works even in the case $m = 4$, but requires a Schreier tree and so does not have such good asymptotic properties.

Algorithm 4.4. TestSet(ω, L)
Define $L_\omega = L$ and for $\omega' \in \Omega$, define $L_{\omega'} = L_\omega^g$ where $\omega^g = \omega'$. This is a valid definition

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

since L is G_ω-invariant.

```
1. Choose $\omega_1 \in L_\omega \setminus \{\omega\}$.
2. Choose $\omega_2 \in L_\omega \cap L_{\omega_1} \setminus \{\omega, \omega_1\}$.
3. Let $X = L_\omega \cap L_{\omega_1} \cap L_{\omega_2}$.
4. If $\#X = 3$ then $B := L_\omega \cap L_{\omega_1} \setminus \{\omega_2\}$
5. Else $B := X$.
6. TestSolutionSubset$(B)$.
```

The proof is similar to that of the previous version, and is omitted here. Note, however, that when $m = 4$, the solution subset B constructed by this algorithm may not actually be a translate of A (since when $m = 4$ and A is a solution subset, $\Omega \setminus A$ is also a solution subset).

5. *Using adjacent-point sets*

We now use the results of the preceding sections to give an algorithm to handle one of the classes of groups introduced at the end of Section 2.

For a positive integer k we will use the term ‘partition of k’ to mean a collection of positive integers, possibly with repetitions, which sum to k. For a finite collection S of finite sets, define the function $p(S) = \{\#x \mid x \in S\}$ to give the partition of $\sum_{x \in S} \#x$ corresponding to the sizes of the elements of S (note that repetitions are not eliminated on the right-hand side).

Recall that $\mathcal{Z}$ is one of the classes of 2-homogeneous groups mentioned in Section 2, containing most of the almost simple 2-transitive groups and a few of the affine 2-transitive groups. We will work with subclasses $\mathcal{Z}_i$ of $\mathcal{Z}$. With each subclass $\mathcal{Z}_i$ we will associate two functions. The function M_i will take a degree m, and give a partition of the number $2m-4$. The function r_i will map the degree m to a natural number. We will choose the classes and functions to satisfy two conditions for each H in $\mathcal{Z}_i$. Let $(H, \Gamma) \in \mathcal{Z}_i$ have degree $m = \#\Gamma$, let $w \in \Gamma^{\{2\}}$ and let S be the collection of H_w-orbits on $\Gamma^{\{2\}}$. Then the following conditions will hold:

(i) $p(\{x \in S \mid x \subseteq \Lambda_\Gamma(w) \setminus \{\{w\}\}\}) = M_i(m)$ (so the partition of $2m-4$ corresponding to the H_w-orbits contained in the adjacent-point set is $M_i(m)$);

(ii) $\#\{L \in \mathcal{P}(S \setminus \{\{w\}\}) \mid p(L) = M_i(m)\} \leqslant r_i(m)$, where $\mathcal{P}$ means ‘power set of’ (so the total number of collections of H_w orbits whose lengths partition $2m-4$ in this way is bounded by $r_i(m)$).

We will also ensure that the functions M_i and r_i are easily computable.

A procedure fulfilling Specification 2.4 for the class $\mathcal{Z}_i$ is as follows.

Algorithm 5.1. TestClass$\mathcal{Z}_i$

```
1. Calculate $m$ such that $\binom{m}{2} = \#\Omega$.
2. If there are no groups in $\mathcal{Z}_i$ with degree $m$ then exit with $b :=$ FALSE.
3. Fix $\omega \in \Omega$ and calculate the collection $S$ of $G_\omega$-orbits on $\Omega$.
4. If $\#\{L \in \mathcal{P}(S \setminus \{\{w\}\}) \mid p(L) = M_i(m)\} \leqslant r_i(m)$ then
5.     For all $L \in \mathcal{P}(S \setminus \{\{w\}\})$ with $p(L) = M_i(m)$ do
6.         TestSet$(\omega, L \cup \{\omega\})$; if $b =$ TRUE then exit (with output $b, \Gamma, \eta$ as returned
               by TestSet).
```

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

```
        End for.
    End if.
    Exit with $b :=$ FALSE.
```

The idea is that if (G, Ω) is the action on pairs of one of the groups in $\mathcal{Z}_i$ then the conditions on the lengths of the G_ω-orbits making up $\Lambda_\Gamma(\omega)$ reduce the possibilities for L; not only that, but we know that this restriction must limit us to at most $r_i(m)$ possibilities for L—if there are more, we do not have to test any of them.

The set S can be formed in $O(sn^2)$ time, using the Schreier generators of G_ω. The remainder of the time taken by this procedure is dominated by the $r_i(m)$ calls to the subroutine TestSet, which has complexity $O(sn)$.

Table 1 uses Theorem 2.2 and the classification of finite 2-transitive groups to classify all the 2-homogeneous groups, giving the class ($\mathcal{A}_o$, $\mathcal{A}_e$, $\mathcal{L}$ or one of the classes $\mathcal{Z}_i$) for each type. It gives details of the orbit lengths for $\Lambda_\Gamma(\omega)$ and, for the groups of class $\mathcal{Z}$, for which TestSet is to be used, the function $r_i(m)$. For most types, the 'Description' column contains a normal subgroup of the group in question, and each line represents several groups that contain the group listed as a normal subgroup.

Note that the lengths of the G_ω-orbits making up $\Lambda_\Gamma(\omega)$ are only given for some of the 2-transitive groups and, where they are, the table actually gives the lengths of the orbits of the two-point stabilizer $G_{\alpha,\beta}$ in the 2-transitive action of the group. To obtain the G_ω-orbit lengths for $\Lambda_\Gamma(\omega)$, replace the initial pair of '1's by a single '1', and then double the lengths of all other orbits. This is because if the $G_{\alpha,\beta}$-orbits on Γ are $\{\alpha\}, \{\beta\}, T_1, \dots, T_k$ then the $G_{\{\alpha,\beta\}}$-orbits that make up $\Lambda_\Gamma(\{\alpha, \beta\})$ are $\{\{\alpha, \beta\}\}$ and the orbits $\{\{\alpha, t\} \mid t \in T_i\} \cup \{\{\beta, t^g\} \mid t \in T_i\}$ where g interchanges α and β, and i runs from 1 to k. Such a g exists because the groups under consideration are all 2-transitive. (Note that g normalizes $G_{\alpha,\beta}$ and so T_i^g is always a $G_{\alpha,\beta}$-orbit.)

For certain affine groups, some orbits are denoted by the letter K. This can be read as '1, 1, divisors of e summing to $q-2$'. These orbits depend on what field automorphisms are present in the group. The reason for the notation will be explained later, following Proposition 8.3.

The table is derived from a similar one in [3], with a couple of corrections and some reorganisation. The functions r_i were calculated by hand with assistance from GAP [8]. Note that, in order to calculate these functions, we need to know the sizes of all the G_ω orbits on Ω, and not just the number of such orbits. In most cases it is relatively straightforward to analyse the orbits of G_ω on $\Omega = \Gamma^{\{2\}}$, to give the results as shown. For type 10, the alternating and symmetric groups, we can take $r_i(m) = 1$ except when $m = 7$. For type 11, $\mathrm{PSL}(d, q)$, there are at most $(q+2)/4$ suborbits of size $2(q-1)$ and at most 4 of size $2q^2(q^{d-2}-1)/(q-1)$. Type 16, class $\mathcal{Z}_9$ (symplectic groups over $\mathbb{F}_2$) is harder. In this case the number of suborbits on pairs is at most nine (for any value of d). The author is unaware of any published version of this result; Appendix A contains a calculation of the lengths of these suborbits, from which it follows that $r_9(m) = 1$.

It will be seen from the table that the worst-case asymptotic size of any of the $r_i(m)$ is $O(m^2)$, and so TestClass$\mathcal{Z}_i$ has asymptotic complexity $O(sn^2)$, since we can also find S in that time. It follows that we can handle the whole class $\mathcal{Z} = \bigcup_{i=1}^{14} \mathcal{Z}_i$ in $O(sn^2)$ time. In fact, r_i is bounded by a constant for all classes except $\mathcal{Z}_6$, $\mathcal{Z}_7$ and $\mathcal{Z}_8$, so if we know the G_ω-orbits in advance we can handle all of $\mathcal{Z}$ except these three subclasses in $O(sn)$ time. In Section 6, we present an $O(sn \log n)$ algorithm to handle the class $\mathcal{Z}_6$. The algorithms for class $\mathcal{L}$ (Section 9) handle the remaining classes, in which the groups have socle $\mathrm{PSL}(2, q)$,

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

Table 1: 2-homogeneous groups

(a) Affine Groups

Type	Cl.	Dimension	Field size	Description	$G_{\alpha,\beta}$-orbits	$r_i(m)$
1	$\mathcal{A}_e$	$d > 2$	$q = 2^e > 2$	$T.\,\mathrm{SL}(d, q)$	K, $q^d - q$	
2	$\mathcal{A}_e$	$d > 2$, even	$q = 2^e > 2$	$T.\,\mathrm{Sp}(d, q)$	K, $q^{d-1} - q$, multiples of q^{d-1}	
3	$\mathcal{A}_e$	$d = 6$	$q = 2^e > 2$	$T.\,\mathrm{G}_2(q)$	K, $q^3 - q$, $q^5 - q^3$, multiples of q^5	
4	$\mathcal{Z}_1$	$d > 2$ $d = 4$	$q = 2$	$T.\,\mathrm{SL}(d, 2)$ $T.\mathrm{A}_7$	1, 1, $2^d - 2$	1
5	$\mathcal{Z}_2$	$d > 2$, even $d = 4$	$q = 2$	$T.\,\mathrm{Sp}(d, 2)$ $T.\mathrm{A}_6$	1, 1, $2^{d-1} - 2$, $2^d - 2^{d-1}$	1
6	$\mathcal{Z}_3$	$d = 6$	$q = 2$	$T.\,\mathrm{G}_2(2)$	1, 1, 6, 24, 32	25
6a	$\mathcal{Z}_4$			$T.\,\mathrm{PSU}(3, 3)$	1, 1, 6, 16, 16, 24	35
7	$\mathcal{A}_e$	$d = 2$	$q = 2^e > 2$	$T.\,\mathrm{SL}(2, q)$	K, multiples of q	
8	$\mathcal{A}_e$	$d = 1$	$q = 2^e > 2$	$G \leqslant \mathrm{A\Gamma L}(1, q)$	K	
9	$\mathcal{A}_o$	$d \geqslant 1$	p odd	$G \leqslant \mathrm{A\Gamma L}(d, p)$		

(b) Almost Simple Groups

Type	Class	Description	Degree (m)	$G_{\alpha,\beta}$-orbits	$r_i(m)$
10	$\mathcal{Z}_5$	A_m, $m > 5$	m	1, 1, $m - 2$	2
11	$\mathcal{Z}_6$	$\mathrm{PSL}(d, q)$, $d \geqslant 3$ A_7 $(d = 4, q = 2)$	$\frac{q^d-1}{q-1}$	1, 1, $q - 1$, $q^2(\frac{q^{d-2}-1}{q-1})$	$q + 2$
12	$\mathcal{Z}_7$ $\mathcal{Z}_8$	$\mathrm{PSL}(2, q)$, $q \geqslant 4$	$q + 1$	1, 1, $q - 1$ or 1, 1, $\frac{1}{2}(q - 1)$, $\frac{1}{2}(q - 1)$	$\frac{1}{2}(q + 2)$ $(q + 2)^2$
13	$\mathcal{L}$	$\mathrm{Sz}(q)$, $q = 2^{2d+1}$, $d \geqslant 1$	$q^2 + 1$		
14	$\mathcal{L}$	$\mathrm{PSU}(3, q)$, $q > 2$	$q^3 + 1$		
15	$\mathcal{L}$	$\mathrm{R}(q)$, $q = 3^{2d+1}$, $d \geqslant 1$ $\mathrm{P\Gamma L}(2, 8)$ $(q = 3)$	$q^3 + 1$		
16	$\mathcal{Z}_9$	$\mathrm{Sp}(2d, 2)$, $d \geqslant 3$	$2^{d-1}(2^d \pm 1)$	1, 1, $2(2^{d-2} \pm 1)(2^{d-1} \mp 1)$, $2^{2(d-1)}$	1
17	$\mathcal{Z}_{10}$	M_k, $k \in \{11, 12, 22, 23, 24\}$	k	1, 1, $k - 2$	1
	$\mathcal{Z}_{11}$	M_{11}	12	1, 1, 10	1
	$\mathcal{Z}_{12}$	$\mathrm{PSL}(2, 11)$	11	1, 1, 3, 6	8
	$\mathcal{Z}_{13}$	HS	176	1, 1, 12, 72, 90	6
	$\mathcal{Z}_{14}$	Co_3	276	1, 1, 112, 162	1

as these groups are of Lie rank one, just as are the other groups of class $\mathcal{L}$. The algorithm presented there runs in nearly linear time, so is asymptotically better than that presented here. The main reason for including these groups in the class $\mathcal{Z}$ is that in practice one may find the the simpler algorithm presented here more useful.

6. *Improved algorithm for* $\mathrm{PSL}(d, q)$

For type 11, $\mathrm{PSL}(d, q)$, we have $M_{11}(m) = \{2(q-1), 2q^2(q^{d-2}-1)/(q-1)\}$ where $m = (q^d - 1)/(q-1)$ and $d \geqslant 3$. There are at most 4 suborbits of the larger size, $2q^2(q^{d-2}-1)/(q-1)$, but there can (depending on what extension of $\mathrm{PSL}(d, q)$ we have) be up to $O(q)$ suborbits of the smaller size, and so the algorithm given in the preceding section has complexity $O(snq)$, which (if $d = 3$) can be as much as $O(sn^{5/4})$. In this section we present a variation on the theme of the preceding section that yields an algorithm for $\mathrm{PSL}(d, q)$ for a fixed pair (d, q) where $m = (q^d-1)/(q-1)$ and $(d, q) \neq (3, 2)$, whose complexity is $O(sn \log n)$. (If $(d, q) = (3, 2)$ then the original method will have to be used.) We use the disparity in the sizes of the two suborbits to avoid searching through all the possibilities for the smaller suborbit.

Algorithm 6.1. TestPSL(d, q)

Input *A set of s generators for a transitive permutation group (G, Ω) of degree $\#\Omega = n = \binom{m}{2}$, a pair of integers (d, q) where $d \geqslant 3$, q is a prime power, $(d, q) \neq (3, 2)$ and $m = (q^d-1)/(q-1)$, a point $\omega \in \Omega$, and a G_ω-orbit $O \subset \Omega$ of size $2(q^{d-2}-1)/(q-1)$.*

Output Either *a G_ω-orbit O' of Ω of size $2(q-1)$ such that if Ω is G-isomorphic to $\Gamma^{\{2\}}$ where (G, Γ) is the 2-transitive action of an extension of* $\mathrm{PSL}(d, q)$ *on* $\mathrm{PG}(d, q)$, *and $O \subset \Lambda_\Gamma(\omega)$ then $\Lambda_\Gamma(\omega) = \{\omega\} \cup O \cup O'$*, or
False, *meaning that it is not the case that Ω is G-isomorphic to $\Gamma^{\{2\}}$ where (G, Γ) is the 2-transitive action of an extension of* $\mathrm{PSL}(d, q)$ *on* $\mathrm{PG}(d, q)$, *and $O \subset \Lambda_\Gamma(\omega)$.*

```
1. Let $L = O \cup \{\omega\}$.
2. Form a Schreier tree of translates of $L$, with edges labelled by the generator of $G$ that
       was used to perform the translation. For each new translate formed, determine
       the size of its intersection with each of the translates already found, and stop
       when a pair $L_1, L_2$ of translates intersect in a set of size greater than 4, or when
       $\lfloor m/2 \rfloor + 1$ translates have been found.
3. If a pair $L_1, L_2$ were found with $\#(L_1 \cap L_2) > 4$ then
4.     Let $g \in G$ be such that $L_1 = L^g$ ($g$ will be obtained from the Schreier tree as a
           word of length $O(m)$ in the generators for $G$).
5.     If there is precisely one $G_{\omega^g}$-orbit $O''$ of size $2(q-1)$ that intersects $L_2 \setminus L_1$
           non-trivially, then return $O' := O''^{g^{-1}}$; else return False.
6. Else
7.     Return False.
8. End if.
9. End.
```

The algorithm runs in $O(sn)$ time.

Proof. Assuming that a solution set A exists with $L \subset \Lambda_A(\omega)$, every translate of L is a subset of the corresponding translate of $\Lambda_A(\omega)$. So the same cases arise as in the proof of correctness of TestSet. Let L_1 and L_2 be distinct translates of L, and Λ_1 and Λ_2 the corresponding translates of $\Lambda_A(\omega)$. It could be the case that $\Lambda_1 = \Lambda_2$, in which case the intersection of L_1 and L_2 must have size at least $1+2q^2(q^{d-2}-1)/(q-1)-2(q-1)$, which is bigger than 4. If (as we supposed) $L_1 \neq L_2$, then L_2 must have non-trivial intersection with the missing suborbit (of the conjugate G_{ω^g} of G_ω) which forms $\Lambda_1 \setminus L_1$, and in fact all of $L_2 \setminus L_1$ must be contained within this suborbit. Thus the suborbit O' found by the algorithm is correct in this case.

Another possibility is that the intersection of Λ_1 and Λ_2 has size m. In this case exactly half of $L_1 \setminus \{\omega_1\}$, and also exactly half of $L_2 \setminus \{\omega_2\}$, will lie in the intersection of L_1 and L_2 (here ω_i is the image of ω under the transformation used to obtain L_i from L). Therefore $L_1 \cap L_2$ has size at least $1 + q^2(q^{d-2} - 1)/(q - 1) - (q - 1)$, which is greater than 4 since $(d, q) \neq (3, 2)$, and at most $1 + q^2(q^{d-2} - 1)/(q - 1)$, and is contained within the intersection $\Lambda_1 \cap \Lambda_2$.

In this case we study the geometry more closely to show that $L_2 \setminus L_1$ will meet exactly one suborbit of size $2(q - 1)$, which will be the one needed to extend L_1 to fill the whole of Λ_1. Let $V = \mathbb{F}_q^d$, with $\Gamma = \mathrm{PG}(V)$. We have $\Lambda_i = \Lambda_A(\{\alpha_i, \beta_i\})$ for some $\alpha_i, \beta_i \in \Gamma$ and $i = 1, 2$. Since $\#(\Lambda_1 \cap \Lambda_2) = m$, we may assume that $\alpha_1 = \alpha_2$, and that the three points α_1, β_1 and β_2 are distinct. Let $a_i, b_i \in V$ be such that $\alpha_i = \langle a_i \rangle$ and $\beta_i = \langle b_i \rangle$ for $i = 1, 2$. Let $X_i = \{x \in \Gamma \setminus \{\alpha_i, \beta_i\} \bigm| x \subset \langle a_i, b_i \rangle\}$ and $Y_i = \{y \in \Gamma \bigm| y \not\subset \langle a_i, b_i \rangle\}$, so the orbits of G_{α_i,β_i} on Γ are $\{\alpha_i\}$, $\{\beta_i\}$, X_i and Y_i, and $\#X_i = q - 1$, $\#Y_i = q^2(q^{d-2} - 1)/(q - 1)$. Finally, we have that $L_i = \{\{\alpha_i, \beta_i\}\} \cup \bigcup_{y\in Y_i} \{\{\alpha_i, y\}, \{\beta_i, y\}\}$.

We consider the $G_{\{\alpha_1,\beta_1\}}$-orbits that meet $L_2 \setminus L_1$. There are two cases, depending on whether β_2 lies in X_1 or in Y_1. If $\beta_2 \in X_1$ then $b_2 \in \langle a_1, b_1 \rangle$, and in fact $\langle a_1, b_1 \rangle = \langle a_2, b_2 \rangle$, so $Y_1 = Y_2$. It follows that any $G_{\{\alpha_1,\beta_1\}}$-orbit containing a pair $\{\beta_2, y\}$ for some $y \in Y_2$ must have size at least that of Y_1, since $y \in Y_1$ and $\beta_2 \notin Y_1$. The only other $G_{\{\alpha_1,\beta_1\}}$-orbit containing a pair in $L_2 \setminus L_1$ is the one that contains $\{\alpha_2, \beta_2\}$, which is the suborbit $\bigcup_{x\in X_1} \{\{\alpha_1, x\}, \{\beta_1, x\}\}$ of size $2(q - 1)$ that equals $\Lambda_1 \setminus L_1$, and is thus the suborbit that we seek. As $\#Y_1 > 2(q - 1)$, this is the only $G_{\{\alpha_1,\beta_1\}}$-orbit of size $2(q - 1)$ that meets $L_2 \setminus L_1$.

On the other hand, if $\beta_2 \in Y_1$ then any $G_{\{\alpha_1,\beta_1\}}$-orbit containing a pair $\{\beta_2, y\}$ for some $y \in Y_2$ must have size at least $\frac{1}{2}\#Y_1$, since $\beta_2 \in Y_1$. The only other pairs in $L_2 \setminus L_1$ are $\{\alpha_1, y\}$ for $y \in Y_2 \setminus (Y_1 \cup \{\beta_1\})$; as $Y_2 \setminus Y_1 = X_1 \cup \{\beta_1\}$, these pairs lie in the $G_{\{\alpha_1,\beta_1\}}$-orbit of length $2(q - 1)$ that we are seeking. As $\frac{1}{2}\#Y_1 > 2(q - 1)$ since $(d, q) \neq (3, 2)$, this is the only $G_{\{\alpha_1,\beta_1\}}$-orbit of size $2(q - 1)$ that meets $L_2 \setminus L_1$.

So we have seen how two of the possibilities for $\Lambda_1 \cap \Lambda_2$ enable us to identify the suborbit that we need to extend L_1 to the whole of Λ_1. The final possibility is that $\Lambda_1 \cap \Lambda_2$ has size 4, in which case $L_1 \cap L_2$ has size at most 4. Observe, as in TestSet, that by the time we have considered $\lfloor m/2 \rfloor + 1$ distinct translates of L, we must have found two of them, L_1 and L_2, say, whose intersection does not fall into this third category.

The algorithm can be implemented in much the same way as TestSet, and in the same time bound. □

Obviously, the intention is that if TestPSL returns a suborbit O' then we form $L = O \cup O' \cup \{\omega\}$ and pass it to TestSet, which also has complexity $O(sn)$ and only needs to be called once after each call to TestPSL.

To complete our discussion of groups of Type 11, we need to consider how many different pairs (d, q) there can be, satisfying $(q^d - 1)/(q - 1) = m$ for a fixed value of m. Note that there can be more than one: *e.g.*, $31 = (2^5 - 1)/(2 - 1) = (5^3 - 1)/(5 - 1)$. It is easy to see that for a particular value of d there can be at most one suitable value of q, as the equation

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

$x^{d-1} + \cdots + x + 1 = m$ has exactly one positive real solution. As $m > q^{d-1} \geqslant 2^{d-1}$, we must have $d - 1 \leqslant \log m$ (logarithm to base 2), so there are $O(\log m)$ possible values of d and so $O(\log m)$ pairs (d, q).

Since there are at most 4 suborbits of size $2(q^{d-2} - 1)/(q - 1)$ and so at most 4 calls to TestPSL are needed for each pair (d, q), we can conclude that at most $O(\log m)$ calls to TestPSL, and accompanying calls to TestSet, are necessary to handle the class $\mathcal{Z}_{11}$, which can therefore be handled in $O(sn \log n)$ time.

7. *Affine groups of odd characteristic*

We consider now the class of affine 2-homogeneous groups, and look for a procedure satisfying Specification 2.4, taking $\mathcal{C}$ to be this class. Thus Γ is a vector space V of dimension r over a prime field $\mathbb{F}_p$, so $m = p^r$. Then G is a subgroup of $\mathrm{AGL}(V) = \mathrm{AGL}(r, q)$ containing the translation subgroup V. In this section we shall assume that $p > 2$, and that (G, Γ) is a group from the class $\mathcal{A}_o$, the class of all 2-homogeneous affine groups defined over finite fields of odd characteristic.

First, we give an important lemma, which applies for both odd and even primes p.

Lemma 7.1. *Let $G = \mathrm{AGL}(r, p)$, and $V = \mathbb{F}_p^r$. Let $x_1, \ldots, x_t$ be a collection of vectors in V, such that the number t of vectors is not divisible by p. Then there exists a vector $v \in V$ such that for all $g \in G$ which leave the collection of x_i invariant,* (i.e., *there exists a permutation π_g of $\{1, \ldots, t\}$ such that $x_i^g = x_{i\pi_g}$ for all i), then $v^g = v$.*

Proof. Let $v = t^{-1} \sum_{i=1}^t x_i$, so v is the average of the vectors x_i, which exists since t and p are coprime. Let $g \in G$ leave the collection of x_i invariant (as described); then

$$\sum_{i=1}^{t} x_i^g = \sum_{i=1}^{t} x_i. \tag{2}$$

Write g as a composite of a linear transformation $h \in G_0$, followed by a translation by a vector $z \in V$, viz. $x^g = x^h + z$ for all $x \in V$. Then, using the linearity of h,

$$\begin{aligned} v^g &= t^{-1}\left(\sum_{i=1}^{t} x_i\right)^h + z \\ &= t^{-1} \sum_{i=1}^{t} x_i^h + z \\ &= t^{-1} \sum_{i=1}^{t} (x_i^g - z) + z. \end{aligned}$$

Then, by (2),

$$v^g = t^{-1} \sum_{i=1}^{t} x_i - t^{-1} t z + z = v.$$

□

We apply this here in the case $p \neq 2$, and later in the case $p = 2$.

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

Corollary 7.2. *Let p be an odd prime, $G \leqslant \mathrm{AGL}(r, p)$ and $V = \mathbb{F}_p^r$. Suppose Δ is a block of imprimitivity of $(G, V^{\{2\}})$ of size not divisible by p. Then there exists $v \in V$ such that the setwise stabilizer $G_{\{\Delta\}}$ of Δ is a subgroup of G_v.*

Proof. Write $\Delta = \{\{u_i, v_i\} \bigm| 1 \leqslant i \leqslant \#\Delta\}$. In Lemma 7.1, take $t = 2\#\Delta$ and $x_{2i-1} = u_i$, $x_{2i} = v_i$. □

In the following theorem we make use of the one-to-one inclusion preserving correspondence between blocks of imprimitivity containing a point ω and subgroups containing the stabilizer G_ω (see [4, Theorem 1.5A], for example).

Theorem 7.3. *Let p be an odd prime, $G \leqslant \mathrm{AGL}(r, p)$ and $V = \mathbb{F}_p^r$, and suppose that G is transitive on $V^{\{2\}}$. Let Δ be a maximal block of $(G, V^{\{2\}})$ subject to $\#\Delta$ not being divisible by p. Then $G_{\{\Delta\}} = G_v$ for some vector $v \in V$, and $\#\Delta = (p^r - 1)/2$, so the block system Δ^G has size p^r.*

Proof. Fix $\omega \in \Delta$, so $G_\omega \leqslant G_{\{\Delta\}}$. By Corollary 7.2, there is $v \in V$ such that $G_{\{\Delta\}} \leqslant G_v$. Then $G_\omega \leqslant G_v$ and so ω^{G_v} is a block of $(G, V^{\{2\}})$ that contains Δ and has size

$$|G_v : G_\omega| = |G : G_\omega|/|G : G_v| = (p^r - 1)/2,$$

since G is transitive on both $V^{\{2\}}$ and V. This is not divisible by p, so by maximality of Δ, we have $\Delta = \omega^{G_v}$, and so $G_{\{\Delta\}} = G_v$. By transitivity, the block system Δ^G has size $\#\Omega/\#\Delta = p^r$. □

Since transitive actions of the same group with identical point stabilizers are isomorphic, Theorem 7.3 suggests the following approach to Specification 2.4 for the class $\mathcal{A}_o$ of 2-homogeneous affine groups defined over fields of odd characteristic.

Algorithm 7.4. AffineFindBlocksOdd

```
1. $b$ := FALSE.
2. If there does not exist an integer $d$ and odd prime $p$ such that $\#\Omega = n = \binom{m}{2}$ where
        $m = p^d$ then Exit.
3. Set $B$ to be a system of blocks of imprimitivity for $(G, \Omega)$ where the blocks have size
        not divisible by $p$, and are maximal in this respect.
4. If $\#B \neq m$ then Exit.
5. Decide whether the group actions $(G, \Omega)$ and $(G, B^{\{2\}})$ are isomorphic, and if so, set
        $b$ := TRUE and find a solution $(\Gamma, \eta)$.
6. Exit.
```

The block system B can be found by using an adaptation of an algorithm of Schönert and Seress [7]. The algorithm given in [7] tests deterministically in $O(n \log^3 \#G + ns \log \#G)$ time whether a transitive group is primitive by finding one minimal block; the authors mention the possibility of extending it to find all minimal blocks, which enables us to find B by making a series of at most $\log n$ calls to this algorithm.

It is claimed in [7] that the extended algorithm runs in the same deterministic, nearly linear, time bound as the original algorithm, which is true if the orbits of G_ω are known in advance, or if an explicit $O(\log \#G)$ bound on the length of subgroup chains in G is known in advance (otherwise every block found has to be checked, which takes $O(sn)$ time, and there can be $O(n)$ blocks to test). In this case we know that if G is a subgroup of the affine group $\mathrm{AGL}(d, p)$ then $\#G \leqslant p^{d(d+1)}$ so $\log \#G \leqslant \frac{1}{4}(\log n + 2) \log n$. As a consequence, we have an explicit bound on the length of a chain of subgroups of G_ω, and so we obtain a deterministic nearly linear time algorithm, even in the case where we do not know the orbits of G_ω in advance.

Similar techniques can be used to improve the calculation of Schreier trees. The cube-doubling routine from [1] normally builds a Schreier tree of depth at most $2\log \#G$ in $O(sn + n\log^2 \#G)$ time, so now the Schreier tree can be assumed to have depth at most $\log^2 n$, and can be built in $O(sn + n\log^4 n)$ time. As before, if this bound for $\log \#G$ is exceeded during the calculation, we can stop, knowing that the group cannot be affine acting on pairs.

(In practical implementations, the Schreier tree would normally be calculated by the straightforward breadth-first search method, which is quick and usually yields a much shallower tree than the $O(n)$ worst case.)

A system of look-up tables will enable us to compute efficiently in B, and thus in $B^{\{2\}}$. To test whether the actions are isomorphic, we first check that $(G, B^{\{2\}})$ is transitive, and then search for a point ω' in $B^{\{2\}}$ that is fixed by G_ω for some point $\omega \in \Omega$. If ω' exists then the map $\omega'^g \mapsto \omega^g$ determines a G-isomorphism. We use a generating set for G_ω to find ω', and the time taken is $O(tn)$ where t is the size of the generating set for G_ω. If no better generating set is available, then the sn Schreier generators can be calculated one by one, and their fixed points in $B^{\{2\}}$ found; this can be done in $O(sn^2)$ time if a whole transversal is explicitly calculated and stored in a preprocessing step.

A better technique would be to calculate a sufficiently large random subset of the Schreier generators, such that the orbits of the subgroup H of G_ω generated by this set do not change with the addition of, say, one or two more random Schreier generators, and choose ω' in $B^{\{2\}}$ to be fixed by this subgroup H. If no such ω' exists, then the two actions cannot be isomorphic; if such an ω' does exist, we can then find the isomorphism η explicitly by doing an orbit calculation on (G, Ω), starting at ω, and, as each new point is discovered, calculating the corresponding point in $B^{\{2\}}$ from those already calculated. This will take $O(sn)$ time.

From this isomorphism we can construct a set which should be (if ω' is in fact invariant under the whole of G_ω) a solution set, and test this set using Algorithm 3.5 (TestSolutionSubset), again in $O(sn)$ time. If this test is failed then we go back and calculate more Schreier generators, thus enlarging the subgroup H of G_ω, and try again. This technique therefore gives a randomized algorithm that cannot give a wrong answer, and although it is not easy to give a useful estimate on the expected number of Schreier generators needed, it will lead to a practical Las Vegas solution, most of which runs in nearly linear time.

It is a theme of the rest of this paper to show that if the orbits of G_ω are known in advance, then there is a deterministic nearly linear algorithm available for Specification 2.4 for the appropriate class $\mathcal{C}$. Although there seems to be no obvious way to construct an isomorphism between Ω and $B^{\{2\}}$ in this time bound if just these orbits are available, we can do so (by the above method) if a suitably small subset of G_ω is available that generates a subgroup that has the same orbits on Ω as G_ω does. In fact, we only need a subgroup of G_ω that has the same fixed points as G_ω, and it is easy to see how this may be constructed from the orbits of G_ω in nearly linear time by iterating over the height of a subgroup chain in such a group. At each stage we choose a point ω' that is not moved by the subgroup so far constructed but is moved by G_ω, and look at the action of the Schreier generators on ω', until one is found that moves it. Of course, there is a certain pointlessness to this argument (how did we obtain the orbits of G_ω in the first place?) but it has been included here for

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

completeness, to show that there is a deterministic nearly linear algorithm for this case if only the orbits of G_ω are known in advance.

Returning to the situation where no extra information about G_ω or its orbits is known in advance, note that if Algorithm 3.5 could be adapted to calculate (in nearly linear time) an element of $G_\omega \setminus H$ when it failed, then we would have a deterministic nearly linear algorithm; alternatively, if a suitably fast technique for producing genuinely random Schreier generators were available, then we would have a nearly linear Las Vegas algorithm. Finally, the Monte-Carlo method in [**1**] could be used to provide a strong generating set with high probability, and this can be used to obtain generators for G_ω. As explained earlier, the $\log \#G$ terms in the analysis of this algorithm can be improved to $\log^2 n$ in this situation using the known bound on the size of the groups being sought. The technique outlined above can then be used to convert this into a true nearly linear Las Vegas algorithm, albeit not in a way that would actually be used in practice unless a strong generating set were needed for other reasons.

Therefore AffineFindBlocksOdd can be implemented deterministically in $O(sn^2)$ time, deterministically in nearly linear time if the orbits of G_ω are known in advance, or in nearly linear Las Vegas time.

This concludes the analysis in the case $p > 2$ (class $\mathcal{A}_o$), so we now consider the case $p = 2$ (class $\mathcal{A}_e$).

8. *Affine groups of characteristic* 2

We consider now the class $\mathcal{A}_e$ which contains most of the affine 2-homogeneous groups defined over fields of characteristic 2, and look for a procedure satisfying Specification 2.4, taking $\mathcal{C}$ to be this class. The precise definition of which groups lie in this class is contained in Table 1; we explain below why the class has been defined in this way.

We assume that (G, Ω) is the action on pairs of a permutation group (G, Γ) that lies in the class $\mathcal{A}_e$. From Table 1, there are well-defined parameters (well-defined because the different rows of Table 1 are disjoint) p, q, d, e such that Γ is a vector space V of dimension d over the field $\mathbb{F}_q$, where p is prime (in fact $p = 2$ in this section) and $q = p^e$, so $m = \#\Gamma = q^d = p^{de}$. Then G will be a subgroup of $\mathrm{A\Gamma L}(V) = \mathrm{A\Gamma L}(d, q)$ containing the translation subgroup V. Therefore G_0 is a semilinear group over $\mathbb{F}_q$. Note that, in writing our algorithm, we cannot assume knowledge of these parameters: since we will have $m = q^d = p^{de}$ we will very quickly be able to identify p and the product de, but some work will be needed to determine d and e (and hence q) from the product de.

Some results, principally Proposition 8.1 and Theorem 8.2, will require only that G is a subgroup of $\mathrm{AGL}(de, p)$, and hence that G_0 is a linear group over the prime field $\mathbb{F}_p$ (observe that $\mathrm{A\Gamma L}(d, q)$ can always be regarded as a subgroup of $\mathrm{A\Gamma L}(de, p)$, and that $\mathrm{A\Gamma L}(de, p) = \mathrm{AGL}(de, p)$ as p is a prime). Later results will require that G be defined as a subgroup of $\mathrm{A\Gamma L}(d, q)$ where q is strictly larger than 2. The reader will observe from Table 1 that in order to enable us to make this restriction, the few groups for which this is not the case have been moved from class $\mathcal{A}_e$ to the class $\mathcal{Z}$. In Section 5 we took advantage of the relatively small number of suborbits that these groups have to give an algorithm for them based on entirely different methods.

Our strategy can be outlined as follows: we first show that if we know the action on *ordered* pairs, we can apply the same techniques as in the odd characteristic case to find a block system on which (if G really is the action on pairs of a group in $\mathcal{A}_e$) G acts as it must on Γ. We then study the suborbits of the actions on pairs of groups in the different rows

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

of Table 1, and give a means of identifying candidates for the stabilizer of an ordered pair (which must have index 2 in the stabilizer of an unordered pair, *i.e.*, the stabilizer of a point in Ω).

Recall that $V^{(2)}$ denotes the set of *ordered* pairs of distinct elements of V.

Proposition 8.1. *Let $G \leqslant \mathrm{AGL}(r, p)$ and $V = \mathbb{F}_p^r$. Suppose Δ is a block of $(G, V^{(2)})$ of size not divisible by p. Then there exists $v \in V$ such that $G_{\{\Delta\}} \leqslant G_v$.*

Proof. Write $\Delta = \{(u_i, v_i) \mid 1 \leqslant i \leqslant \#\Delta\}$. In Lemma 7.1, take $t = \#\Delta$ and $x_i = u_i$ for $i = 1, \dots, t$. □

Theorem 8.2. *Let $G \leqslant \mathrm{AGL}(r, p)$ and $V = \mathbb{F}_p^r$, and suppose that G is transitive on $V^{(2)}$. Let Δ be a maximal block of $(G, V^{(2)})$ subject to $\#\Delta$ not being divisible by p. Then the setwise stabilizer $G_{\{\Delta\}}$ of Δ is G_v for some vector $v \in V$, and $\#\Delta = p^r - 1$, so the block system Δ^G has size p^r.*

Proof. This follows from Proposition 8.1 in much the same way that Theorem 7.3 follows from Corollary 7.2. □

Recall from Lemma 2.1 that if (G, Γ) is 2-homogeneous and $\#G$ is divisible by 2 then (G, Γ) is 2-transitive. It follows from Theorem 8.2 that if we can find the action on ordered pairs, then the techniques applied in the odd characteristic case will yield a solution for the class of affine groups defined over fields of characteristic 2 as well.

As G is 2-homogeneous and $\#G$ is divisible by 2, the stabilizer $G_{\alpha,\beta}$ of an ordered pair is always a subgroup of index 2 in the stabilizer $G_{\{\alpha,\beta\}}$ of the corresponding unordered pair. So the problem of finding possibilities for the action on ordered pairs, given the action on unordered pairs, reduces to that of finding relevant subgroups of index 2 in G_ω.

We now make use of the fact that groups in $\mathcal{A}_e$ are defined over fields *larger* than the prime field $\mathbb{F}_2$. Having fixed $\omega \in \Omega$, we find a G_ω-orbit $\Delta \subseteq \Omega$ such that the stabilizer in G_ω of any point in Δ must be contained in the ordered-pair stabilizer in which we are interested. This will enable us to construct the action of G on cosets of $G_{\alpha,\beta}$.

Proposition 8.3. *Let G be a subgroup of $\mathrm{A\Gamma L}(d, q)$ containing the translation subgroup, where $q = 2^e$. Let $V = \mathbb{F}_q^d$, with G acting naturally on V. Let $\omega = \{v_1, v_2\} \in V^{\{2\}}$, and let $H = G_{v_1,v_2}$. Then H is a direct factor of G_ω of index 2. Also, there is a G_ω-invariant subset K of V of size q such that the group induced by H on $K^{\{2\}}$ is cyclic of order f where $f \mid e$. Furthermore, if $e > 1$ then $K^{\{2\}}$ contains a G_ω-orbit of size $2f$.*

This is the reason for the 'K' notation in Table 1.

Proof. Without loss, $\omega = \{0, v\}$. Then $H = G_{0,v}$, and in fact $G_\omega = H \times \langle t_v \rangle$ since $\langle t_v \rangle \trianglelefteq G_\omega$, and t_v (the translation map $x \mapsto x + v$ on V) has order 2. Define $K = \{\alpha v \mid \alpha \in \mathbb{F}_q\}$. Then K is invariant under H and under t_v (as $(\alpha v)^{t_v} = (\alpha + 1)v$), so it is invariant under G_ω. Let $N = H \cap \mathrm{GL}(d, q)$. Then N acts trivially on K since it acts linearly on V and fixes a generator, v, of the one-dimensional subspace K of V. Therefore, N lies in the kernel of the action induced by H on $K^{\{2\}}$.

We know that $\Gamma\mathrm{L}(d, q)$ is the semidirect product of $\mathrm{GL}(d, q)$ by a cyclic group of order e generated by the Frobenius automorphism σ, so H/N is cyclic of order dividing e. Thus H induces a cyclic group C on K (and hence on $K^{\{2\}}$) of order f, say, where $f \mid e$, and C acts on K as a group of field automorphisms. Let α generate the multiplicative group of $\mathbb{F}_q$,

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

and let $g \in C$. Then $(\alpha^r)^g = (\alpha^g)^r$ for any integer r, and so if C does not act faithfully on the orbit X containing αv, then it cannot act faithfully on K, which would contradict the choice of C. As C is abelian and acts faithfully and transitively on X, it must act regularly on X. Therefore X is an orbit of C on K of size f. The required orbit of G_ω acting on pairs is now $\{\{z, x\} \mid z \in \{0, v\},\ x \in X\}$, as this has size $2f$ except when $f = e = 1$. □

If (G, V) lies in $\mathcal{A}_e$, and so can be defined over a field of size 2^e where $e > 1$, and if we know $K^{\{2\}}$, we can use this proposition to find a set of at most three subgroups of index 2 in G_ω, one of which is the desired subgroup H. To do this, we first take a G_ω-orbit Y in $K^{\{2\}}$ of maximum size; this size should be even, and we set $f = \#Y/2$. By the above, G_ω should induce a regular group isomorphic to $C_f \times C_2$ on Y. The desired subgroup (if it exists) will be the pre-image in G_ω of one of the cyclic subgroups of size f of this group. (The number of such cyclic subgroups depends on f: it is one if f is odd, two if f is divisible by 4, and three if $f \equiv 2 \bmod 4$.) The orbit systems of these cyclic subgroups are systems of imprimitivity of G_ω, and can be found quickly (or shown not to exist) using the regularity and cyclicity conditions; the desired subgroup is the kernel of the action of G_ω on one of these block systems. We summarise this as an explicit procedure to find suitable block systems or show that they do not exist.

1. Let Y be a G_ω-orbit in $K^{\{2\}}$ of maximum size.
2. If $\#Y$ is not divisible by 2 then Exit.
3. Else Set $f := \#Y/2$.
4. End if.
5. Fix $y \in Y$, and find a transversal T for $G_{\omega,y}$ in G_ω. (N.B. If (as hoped) Y^{G_ω} is regular, then the images of T are the elements of Y^{G_ω}.)
6. Let Z be the set of orbits $y^{\langle t\rangle}$ as t ranges over those elements of T that induce elements of order f on Y.
7. Check that the number of sets in Z is correct (one if f is odd, two if f is divisible by 4, and three if $f \equiv 2 \bmod 4$).
8. Each set in Z will be a block for G_ω, and the corresponding block systems are the ones required.

Much of the calculation can be carried out in the group induced by G_ω on Y, which is of degree $f = O(\log q)$. If a suitable regularity test is available, it might be worthwhile to check that this group is regular before calculating the orbits in Z.

The small degree of the group (G_ω, Y) means that this calculation is cheap: if s generators are given for G_ω then the above procedure has complexity $O((s + f)f)$, as $O(sf)$ time is required to form the induced action, and subsequent calculations can all be done in $O(f^2)$ time. If we use Schreier generators for G_ω, the complexity of the procedure becomes $O(sn \log n)$, where the timing is now dominated by the calculation of the induced action.

Having found a G_ω-block X (here X is one of the elements of Z in the preceding paragraph) whose stabilizer in G_ω is a candidate for H, the stabilizer of the ordered pair corresponding to the (as yet hypothetical) unordered pair ω, it is now an easy matter to construct the action of G on right cosets of H in G as an action on the cartesian product of Ω and a set $C = \{\textsc{True}, \textsc{False}\}$ of size 2. Fix a transversal T for $G_\omega \in G$. For a generator g_i of G, and a point (ω', x), the image $(\omega', x)^{g_i}$ is given by

$$(\omega', x)^{g_i} = \begin{cases} (\omega'^g, x), & \text{if } X^{t(\omega')gt(\omega'^g)^{-1}} = X \\ (\omega'^g, \neg x), & \text{if } X^{t(\omega')gt(\omega'^g)^{-1}} \neq X \end{cases}$$

where $t(\omega')$ is the element of T mapping ω to ω'. Perform this calculation for all generators and points to get the images of the generators in the new group. The transversal T can be stored in a Schreier tree if desired; testing whether $X^{t(\omega')gt(\omega'^g)^{-1}} = X$ can be done by examining just one point of X, since the translating element lies in G_ω and X is a block for G_ω. Using a look-up table to check membership of X, the whole procedure can be completed in $O(sn)$ time if T is stored explicitly, or $O(snl)$ time if T is stored in a Schreier tree of depth l.

This action can now be tested as described earlier: we find a maximal block in it of size not divisible by 2, and if the corresponding block system has the right size, we run an isomorphism test to see if we have a solution. If not, we repeat for each $X \in Z$, and if no solution is found over all these tests, or if we failed when trying to calculate X, we may conclude that no solution exists for groups in the class $\mathcal{A}_e$ of affine groups defined over fields of characteristic 2 where the field is larger than $\mathbb{F}_2$, and where $K^{\{2\}}$ is known.

We now turn to the problem of identifying $K^{\{2\}}$.

Lemma 8.4. *Let x, e be positive integers and $q = 2^e$. Then*

(i) *if $e > 1$ and $\frac{1}{2}q$ divides $4e$ then $2 \leqslant e \leqslant 4$;*

(ii) *if $e > 1$ and $x > 2$ then $q^x - q > 4e(x+3)$;*

(iii) *if $e = 1$ and $x > 4$ then $q^x - q > 4x + 4$.*

Proof. The proof is straightforward, and is omitted here. □

Fix a point $\omega \in \Omega$. Given a positive integer e, define the subset Q_e of Ω by

$$Q_e = \{\omega' \in \Omega \mid \#(\omega'^{G_\omega}) \text{ divides } 2e\},$$

that is, the union of all G_ω-orbits of size dividing $2e$.

Proposition 8.5. *Suppose (G, V) is a group of one of the types 1 to 8 in Table 1, and let $\Omega = V^{\{2\}}$. Let q, e and d be the corresponding parameters from this table, and let $r = de$, so $m = \#V = 2^r$. Then one of the following holds:*

(i) *the set Q_e has size $2^{e-1}(2^e - 1)$;*

(ii) $G = \mathrm{ASL}(2,4)$;

(iii) $G = \mathrm{A\Sigma L}(2,4) = \mathrm{ASL}(2,4):2$*; or*

(iv) $\mathrm{ASL}(2,16) \leqslant G \leqslant \mathrm{A\Sigma L}(2,16) = \mathrm{ASL}(2,16):4$.

Furthermore, in the first case, e is the largest integer dividing r such that Q_e has size $2^{e-1}(2^e - 1)$. In the second case, the largest integer e^* dividing r such that Q_{e^*} has size $2^{e^*-1}(2^{e^*} - 1)$ is 4. In the third and fourth cases there is no such integer e^*.

Proof. Define e^* to be the largest integer dividing r such that Q_{e^*} has size $2^{e^*-1}(2^{e^*} - 1)$, or 0 if no such integer exists. We suppose that $e^* \neq e$, and show that then either $e^* = 0$ and one of the third and fourth cases arises, or that $e^* = 4$, and $G = \mathrm{ASL}(2,4)$.

From Proposition 8.3, we know that Q_e contains $K^{\{2\}}$. Therefore (under our assumption that $e^* \neq e$), there is an orbit T of G_ω on Ω of size dividing $2r$, and which is not contained in $K^{\{2\}}$, since either $e^* > e$, or Q_e strictly contains $K^{\{2\}}$. Clearly, therefore, it is impossible for G to be of type 8 under our assumption, as for that type $K = V$. If U is an orbit of $G_{\alpha,\beta}$ on V such that pairs in T contain elements of U, then $\frac{1}{2}\#U \mid \#T$, and $\#T \mid 2e^* \mid 2r$ by hypothesis. It follows that if U is an orbit of G_ω (or $G_{\alpha,\beta}$) on V, and U is of minimal size subject to not being contained in K, then $\frac{1}{2}\#U \leqslant 2r$.

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

From the lists of $G_{\alpha,\beta}$ orbit lengths in Table 1, and Lemma 8.4, it is immediate that if G is not of type 7, then $e = 1$ and d is 3, 4 or (for types 6, 6a) 6. It is straightforward to check manually or by GAP [8] that in all these remaining cases, $e^* = e = 1$. Similarly, it is immediate from Table 1 and Lemma 8.4 that if G is of type 7 (so $d = 2$) then $e = 2$, 3 or 4; again, it is straightforward to check that in this case if $e = 3$ then also $e^* = 3$.

The suborbit lengths of the groups containing $\mathrm{ASL}(2, 4)$ and of some of those containing $\mathrm{ASL}(2, 16)$ are displayed in Table 2, with the value of e^*. The table was computed using GAP [8].

Table 2: Suborbit lengths of exceptions in Proposition 8.5

Group	Number of G_ω-orbits on $\Omega = V^{\{2\}}$								e^*
Lengths:	1	2	4	6	8	12	16	24	
$\mathrm{ASL}(2, 4)$	2	11			12				4
$\mathrm{ASL}(2, 4) : 2$	2	1	5		4		4		0
$\mathrm{AGL}(2, 4)$	2	2		3				4	2
$\mathrm{A\Gamma L}(2, 4)$	2		1	1		1		4	2
Lengths:	1	2	4	8	16	32	64	128	
$\mathrm{ASL}(2, 16)$	8	56		225		960			0
$\mathrm{ASL}(2, 16) : 2$	4	10	24	9	108	48	456		0
$\mathrm{ASL}(2, 16) : 4$	2	1	5	13	4	58	22	228	0
Lengths:	1	2	24	40	96	160			
$\mathrm{ASL}(2, 16) : 3$	8	56	75		320				4
$\mathrm{ASL}(2, 16) : 5$	8	56		45		192			4

The table includes groups isomorphic to each of the three subgroups of $\mathrm{A\Gamma L}(2, 16)$ which contain $\mathrm{ASL}(2, 16)$ as a maximal proper subgroup. It also includes groups isomorphic to each subgroup of $\mathrm{A\Gamma L}(2, 16)$ in which the index of $\mathrm{ASL}(2, 16)$ divides 4. Suppose $G \leqslant \mathrm{A\Gamma L}(2, 16)$, but G is not isomorphic to a group in the table. Then (as $|\,\mathrm{A\Gamma L}(2, 16) : \mathrm{ASL}(2, 16)| = 60$), $|G : \mathrm{ASL}(2, 16)|$ must be divisible by 3 or 5, and so G must contain a subgroup H isomorphic to one of the groups in the last two rows of the table. The group H has a correct e^* of 4, and since the suborbits of G are unions of suborbits of H, yet K remains the same, this means that e^* for G also has the correct value, 4.

The claimed results now follow from this discussion and the entries in the table. □

Let $\mathcal{A}_x$ be the class of groups consisting only of $\mathrm{ASL}(2, 4)$, $\mathrm{ASL}(2, 4) : 2$, and those groups G with $\mathrm{ASL}(2, 16) \leqslant G \leqslant \mathrm{ASL}(2, 16) : 4$. Let $\mathcal{A}_e^* = \mathcal{A}_e \setminus \mathcal{A}_x$. We now have a procedure for the class $\mathcal{A}_e^*$.

1. If there does not exist an integer r such that $\#\Omega = n = \binom{m}{2}$ where $m = 2^r$, then conclude that (G, Ω) is not a group from $\mathcal{A}_e^*$ acting on pairs.
2. Fix $\omega \in \Omega$ and calculate the G_ω-orbits on Ω.
3. Calculate the largest integer e^* dividing r and such that Q_{e^*} has size $2^{e^*-1}(2^{e^*} - 1)$.
4. If $e^* > 1$ then apply previously described techniques, assuming that $K^{\{2\}} = Q_{e^*}$ (and therefore that $e = e^*$ and $d = r/e$).

5. Else conclude that (G, Ω) is not a group from $\mathcal{A}_e^*$ acting on pairs.
6. End if.

Observe that if (G, Ω) is one of the affine groups in $\mathcal{Z}$ (*i.e.*, an affine group defined over $\mathbb{F}_2$) acting on pairs, then the above test will give $e^* = 1$. This may be used to speed up an implementation of the procedure for groups of class $\mathcal{Z}$.

The class $\mathcal{A}_x$, which contains only five groups, need not detain us long. The simplest technique here is to divide it into five classes, each containing one group. The first test in each case will be to check that there are the correct number of suborbits of each length, by comparison with Table 2. If that test is passed, we test a series of candidates for a largest G_ω-orbit lying in $K^{\{2\}}$, using information about the G_ω-orbit lengths in $K^{\{2\}}$. For example, for $\mathrm{ASL}(2, 16) : 4$, these orbits have length 8, and we test two of the G_ω-orbits of this length, since twelve out of the thirteen G_ω-orbits of length 8 lie in $K^{\{2\}}$. In the other cases the number of orbits to be tested is between one and ten. Each candidate G_ω-orbit is tested in the same way that a largest G_ω orbit in $K^{\{2\}}$ is tested in the procedure for $\mathcal{A}_e^*$.

As in the odd characteristic case, this algorithm can be implemented deterministically in $O(sn^2)$ time, or in nearly linear Monte-Carlo time. The analysis is the same as for that case, except that we have to find relevant subgroups of G_ω of index 2, which takes $O(sn \log n)$ time, and that we get up to three actions of degree $2n$, instead of just one of degree n, to which to apply the block-finding routine.

9. *The remaining almost simple groups*

The class $\mathcal{L}$ consists of the groups (G, Γ) that contain a normal subgroup S isomorphic to one of the simple groups $\mathrm{PSU}(3, q)$ (for q a prime power greater than 2) (acting on isotropic points), $\mathrm{Sz}(q)$ (for q an odd power of 2 greater than 2) or $\mathrm{R}(q)$ (for q an odd power of 3 greater than 3) as a normal subgroup, and the degree m is $q^3 + 1$, $q^2 + 1$ or $q^3 + 1$ respectively. We also include the non-simple Ree group, $\mathrm{R}(3) \cong \mathrm{P\Gamma L}(2, 8)$ acting on 28 points in the class $\mathcal{L}$. For the moment, however, we shall ignore it; it will be discussed at the end of this section.

The groups in $\mathcal{L}$ are groups of Lie type, and all have Lie rank one; the other family of groups of Lie rank one is $\mathrm{PSL}(2, q)$, for q a prime power greater than 3, acting on projective points, which has many similarities to these families, but sufficient differences for us to treat it in the class $\mathcal{Z}$. However the results of this section apply to this family of groups as well. In what follows, $\mathcal{L}'$ will be used to denote the class of 2-transitive groups that either lie in $\mathcal{L}$, or whose socle is $\mathrm{PSL}(2, q)$ for some prime power $q > 3$.

We refer the reader to [**4**, pp. 248–252] for a brief description of these groups. The notation here follows that of [**4**] in most respects.

We consider the groups S, which are themselves all 2-transitive. The two-point stabilizer $S_{\alpha,\beta} = H$ is cyclic of order $(q^2 - 1)/d$ for $\mathrm{PSU}(3, q)$, where d is $(3, q + 1)$, and cyclic of order $q - 1$ for the other cases. The stabilizer S_α is the semidirect product of a Sylow p-subgroup T of S by H (where p is the prime dividing q), and T acts regularly on $\Gamma \setminus \{\alpha\}$ so $\#\Gamma = \#T + 1$.

Proposition 9.1. *Let (G, Γ) be a member of the class $\mathcal{L}'$. Let $\alpha, \beta \in \Gamma$. Let X be a subgroup of G, with $G_{\alpha,\beta} \leqslant X$. Suppose $|X : G_{\alpha,\beta}|$ is a power of the unique prime p dividing $m - 1$. Then either X is contained in G_α or G_β, or $p = 2$ and $X = G_{\{\alpha,\beta\}}$.*

Proof. We shall assume that $X \nleqslant G_\alpha$, $X \nleqslant G_\beta$ and $X \neq G_{\{\alpha,\beta\}}$, and seek a contradiction. Let $c = |X : G_{\alpha,\beta}|$, so certainly $c > 1$. Let $Y = X \cap S$ where as before S is the socle of G.

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

Since S is 2-transitive, $G = G_{\alpha,\beta}S$. Therefore $G = XS$ and by the second isomorphism theorem, $|X : Y| = |XS : S|$, and so $|X : Y| = |G_{\alpha,\beta} : H|$ and therefore $|Y : H| = c$. By the modular law, $G_{\alpha,\beta}(S \cap X) = G_{\alpha,\beta}S \cap X$, since $X \geqslant G_{\alpha,\beta}$, so since $G = G_{\alpha,\beta}S$ we have $G_{\alpha,\beta}Y = X$. Therefore $Y \nleqslant S_\alpha$, $Y \nleqslant S_\beta$, and $Y \neq S_{\{\alpha,\beta\}}$. Since $N_S(H) = S_{\{\alpha,\beta\}}$, it follows that H is not normal in Y.

Recall that S_α is the extension of a Sylow p-subgroup T (of size $m-1$) of S by H. This means that $N_S(T) = S_\alpha$, and so there is precisely one Sylow p-subgroup of S fixing each point of Γ; as T acts regularly on $\Gamma \setminus \{\alpha\}$, the intersection of any two Sylow p-subgroups of S is trivial.

Let P be a Sylow p-subgroup of Y, so $\#P = c$ since $\#H$ is not divisible by p. We also get $P \cap H = 1$ since the orders are coprime, and therefore $Y = PH$. By Sylow's Theorem, P is contained in a Sylow p-subgroup of S, and so P has a unique fixed point γ and acts semi-regularly on $\Gamma \setminus \{\gamma\}$. Since $Y \nleqslant S_\alpha$ and $Y \nleqslant S_\beta$, $\gamma \notin \{\alpha, \beta\}$. As $P \leqslant Y_\gamma$ and $Y = PH$, we get that $Y = Y_\gamma H$ and so H is transitive on the Y-orbit γ^Y. That is, γ^Y is an H-orbit on $\Gamma \setminus \{\alpha, \beta\}$. The possible lengths of the H-orbits on $\Gamma \setminus \{\alpha, \beta\}$ are $q-1$, $(q-1)/2$, q^2-1 and $(q^2-1)/3$ (see, for example, [3, Table 1]).

We also know that $\#\gamma^Y \equiv 1 \bmod c$, since P acts semi-regularly on $\gamma^Y \setminus \{\gamma\}$. Now c and q are both powers of the prime p. If $q-1 \equiv 1 \bmod c$, or if $q^2-1 \equiv 1 \bmod c$, then $p = c = 2$. If $(q-1)/2 \equiv 1 \bmod c$, then $p = c = 3$. If $(q^2-1)/3 \equiv 1 \bmod c$, then $p = 2$ and c is 2 or 4. Thus $2 \leqslant c \leqslant 4$. However c cannot be 2 as that would imply that $H \lhd Y$, which we know contradicts our hypothesis.

Therefore either (case 1) $p = c = 3$ and $S = \mathrm{R}(q)$ (because $\mathrm{R}(q)$ is the only possibility with an H-orbit of length $(q-1)/2$, or (case 2) $p = 2$, $c = 4$ and $S = \mathrm{PSU}(3, q)$ where $3 \mid q+1$ (as this is the only possibility with an H-orbit of length $(q^2-1)/3$).

Now $\#(\alpha, \beta)^Y = c$ and by hypothesis Y does not fix either α or β. Thus in case 1 we must have $\#\alpha^Y = \#\beta^Y = 3$. However, the Y-orbits are unions of H-orbits, and the H-orbits all have length $q-1$ except $\{\alpha\}$, $\{\beta\}$ and one orbit of length $(q-1)/2$. Since q is at least 27, this clearly leads to a contradiction. Similarly for case 2: $\#\alpha^Y$ and $\#\beta^Y$ must each be either 2 or 4. The H-orbits all have length $(q-1)/3$ except for $\{\alpha\}$, $\{\beta\}$ and one orbit of length $q-1$. Since the situation $\alpha^Y = \beta^Y = \{\alpha, \beta\}$ gives $\#(\alpha, \beta)^Y = 2$, rather than 4, this cannot happen unless $(q^2-1)/3$ is 1, 2 or 3, all of which contradict the choice of q.

Thus we have reached a contradiction, and so conclude that either X is contained within G_α or G_β, or $p = 2$ and $X = G_{\{\alpha,\beta\}}$. □

Recall that $\Gamma^{(2)}$ is the set of ordered pairs of distinct elements of Γ.

Corollary 9.2. *Let (G, Γ) be a member of the class $\mathcal{L}'$. Fix $\alpha, \beta \in \Gamma$. Then $\{(\alpha, x) \mid x \in \Gamma \setminus \{\alpha\}\}$ and $\{(x, \beta) \mid x \in \Gamma \setminus \{\beta\}\}$ are blocks of imprimitivity in the action of G on $\Gamma^{(2)}$; they are maximal among blocks in this action containing (α, β) and of size dividing $m-1$; furthermore they are the only such maximal blocks in this action apart from (in the case where $2 \mid m-1$) the block $\{(\alpha, \beta), (\beta, \alpha)\}$.*

Proof. This is an application of the one-to-one correspondence between blocks of imprimitivity containing (α, β) and subgroups containing $G_{\alpha,\beta}$ (see [4, Theorem 1.5A], for example). Maximality of the blocks given is a consequence merely of their size; the remainder of the result follows from Proposition 9.1. □

Observation 9.3. *Suppose D is one of the two blocks of imprimitivity of $\Gamma^{(2)}$ in Corollary 9.2. Then the subset $A = \{\{x, y\} \mid (x, y) \in D\}$ of $\Gamma^{\{2\}}$ is a solution subset.*

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

We can perform the projection from $\Gamma^{(2)}$ to $\Gamma^{\{2\}}$ even if we only know the stabilizer of a point in $\Gamma^{(2)}$ as a subgroup of index 2 in the stabilizer of a point in $\Gamma^{\{2\}}$, and do not know the structure of either set as a set of pairs. In particular, it is straightforward to perform this projection if we have constructed the supposed action on $\Gamma^{(2)}$ from a subgroup of index 2 of G_ω, in the manner indicated in Section 8.

This means that given a subgroup J of index 2 in G_ω, we can solve Specification 2.4 for the class $\mathcal{L}$ under the added condition that $J = G_{\alpha,\beta}$ where $\omega = \{\alpha, \beta\}$ as follows:

```
Form the action on the set $\Delta$ of cosets of $J$ in $G$.
Find a block of imprimitivity $D$ that is of size dividing $m-1$, contains the trivial
        coset $J$ and does not contain the coset $G_\omega \setminus J$, and is maximal among such
        blocks.
If $\#D = m-1$ then
    Project $D$ from $\Delta$ back onto $\Omega$ (by the map $Jx \mapsto \omega^x$), and test to see whether
            the resulting set is a solution subset.
    If we find a solution subset, return $b :=$ TRUE and $\Gamma$, $\eta$ calculated from the
            solution subset. Otherwise, return $b :=$ FALSE.
Else
    Return $b :=$ FALSE;
End if;
End.
```

Obviously, if we have several candidates J for $G_{\alpha,\beta}$ we can test each one in turn by the above procedure.

The subgroup $G_{\alpha,\beta}$ is a subgroup of index 2 in the stabilizer G_ω of a point in the input action. We can find the subgroups of index 2 using the following result, whose proof can be found in Appendix B. Recall that a base for a permutation group is a subset of the domain such that the point-wise stabilizer of the subset is trivial. (Many authors regard a base as an ordered set; however, that will not be necessary for present purposes.)

Theorem 9.4. *Let* (G, Γ) *lie in* $\mathcal{L}'$. *Let* $\omega \in \Omega = \Gamma^{\{2\}}$ *and let* R *be a largest* G_ω*-orbit. Let* $\lambda \in R$. *Then* $G_{\omega,\lambda}$ *has size at most* 2, G_ω *acts faithfully on* R *and either* $R \cup \{\omega\}$ *is the adjacent-point set* $\Lambda_\Gamma(\omega)$ *or there are at most* $2\log_p(m-1)$ *elements* λ' *of* R *such that* $\{\omega, \lambda, \lambda'\}$ *is not a base for* G (*here* p *is the prime dividing* $m-1$).

Proof. This is reasonably straightforward to show; it is proved in Appendix B. □

Note that for groups in $\mathcal{L}$, and indeed some of the groups whose socle is $\mathrm{PSL}(2, q)$, this can be strengthened to the effect that $G_{\omega,\lambda}$ is actually trivial. This also follows from results in Appendix B.

This result enables us to find the action of G_ω on R, and find a subset of R of size at most 2 which is a base for G_ω. We can then enumerate the elements of G_ω by traversing a Cayley graph, identifying the vertices of the graph with the images of the base. This enables us to form the subgroup G_ω^2 as a block in the regular action of G_ω. The index of G_ω^2 in G_ω is 2, 4 or 8, since G_ω is an extension of a cyclic group H by a group of order 2, and then by a cyclic group of field automorphisms. We can factor out G_ω^2 in our regular action of G_ω, inducing an elementary abelian action of G_ω on 2, 4 or 8 points. The 1, 3 or 7 subgroups of index 2 in this induced group yield the subgroups of G_ω of index 2. As there are at most 7 of these, using the previously described algorithm to check all possible subgroups of index 2 is feasible.

We now give an explicit algorithm to perform these tasks.

1. If $m-1$ is not the power of a prime p then Exit.
2. Fix $\omega \in \Omega$, and set S to be the set of G_ω-orbits on Ω.
3. Set R to be an element of S of maximum size.
4. Test whether the set $R \cup \{\omega\}$ is an adjacent-point set $\Lambda_\Gamma(\omega)$ using TestSet.
5. Form the group induced by G_ω on R.
6. Fix $\lambda \in R$. Calculate $G_{\omega,\lambda}$ as a group of permutations of R; if it has size larger than 2 then Exit, else if it is non-trivial, find $\lambda' \in R$ such that $b = \{\lambda, \lambda'\}$ is a base for G_ω; otherwise take $b = \{\lambda\}$ as a base for G_ω.
7. Use the known base for G_ω to enumerate all the elements of G_ω as images of the base b and as permutations of R.
8. Calculate the subgroup M of G_ω generated by the elements g^2 as g runs over G_ω, as a subset containing b of the set of images of b; this is a block in the action of G_ω on base-images.
9. If $x = \#G_\omega/\#M \in \{2, 4, 8\}$ and there are $x-1$ blocks of size $\#G_\omega/2$ that contain M in the action of G_ω on base-images, then return the block systems corresponding to these blocks.
10. Else conclude that (G, Ω) is not the action on pairs of a group in $\mathcal{L}'$.
11. End if.

We form the action of G_ω on R as follows. Fix $\lambda \in R$ and choose $\lambda' \in R \setminus \{\lambda\}$; we will assume that $\{\lambda, \lambda'\}$ is a base for G_ω. By Theorem 9.4 we may have to run the algorithm $O(\log m)$ times with different values for λ' before we are certain that if the group is the action on pairs of a group from $\mathcal{L}'$ then we have at some point considered a base. As $\#R = O(m \log m)$, there are $O(m^2 \log^2 m)$ possible images of the 2-element base, so in $O(m^2 \log^2 m)$ space and $O(sln)$ time (where l is the depth of a Schreier tree) we can evaluate each of the $O(sn)$ Schreier generators for G_ω on the points of the (assumed) base and decide if there are more than $2\#R$ distinct base images. (We use the $O((\#R)^2)$ space to store a table of flags that enables us to decide whether we have seen a particular base image before in constant time.) If that is the case then G_ω is too large, and we may terminate the procedure. Otherwise we can, in the same time bound, and still assuming that $\{\lambda, \lambda'\}$ is a base for G_ω, evaluate on the whole of R the Schreier generators that led to distinct base images, and arrive at a set of $O(\#R)$ generators for the action of G_ω on R. We test whether these elements generate a group that is transitive on R; if not, then $\{\lambda, \lambda'\}$ cannot be a base, and we must return to the start and choose a different λ. However, the converse does not hold: if the group generated is transitive on R we cannot deduce that $\{\lambda, \lambda'\}$ is necessarily a base.

We now calculate $G_{\omega,\lambda}$; here we have $O((\#R)^2)$ Schreier generators, and still assume that $\{\lambda, \lambda'\}$ is a base. There should be at most 2 distinct base images (including the trivial one); if any more are found, we can terminate the whole procedure (as even if $\{\lambda, \lambda'\}$ is not a base, we now know that G_ω is too large). If 2 elements of $G_{\omega,\lambda}$ are found, then we may conclude that either $\{\lambda, \lambda'\}$ is a base for G_ω or (G, Ω) is not the action on pairs of a group of $\mathcal{L}'$, and so we can continue to the next stage of the algorithm. However, if it appears that G_ω acts regularly on R, then we must repeat the whole procedure with other values of λ', and only conclude that that is really the case when we have tried $1 + 2\log_p(m-1)$ different values for λ'.

Thus in $O(sln \log n)$ time and using $O(n \log^2 n)$ space we can obtain: a subset of $\mathrm{Sym}(R)$ of size at most $2\#R$, which generates G_ω if (G, Ω) is the action on pairs of a group in $\mathcal{L}'$;

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

a permutation of R, of order 1 or 2, which (under the same condition) generates $G_{\omega,\lambda}$; and a subset b of R, of size at most 2 and containing λ, which (under the same condition) is a base for G_ω. It is now straightforward to use these to enumerate the elements of G_ω, both as images of b and as permutations of R, in time $O((\#R)^2)$. Then the required block M in the action of G_ω on images of b can be formed in the same time, and the blocks corresponding to subgroups of index 2 in G_ω can be found quickly.

Note that this procedure can be significantly simplified if it is only desired that we check for members of $\mathcal{L}'$ for which G_ω acts regularly on a largest suborbit R; this includes all members of $\mathcal{L}$. In particular, in that case we can assume that $\{\lambda\}$ is a base for G_ω, since if it is not, then the group is not one we are interested in.

The general case would probably be implemented slightly differently from the outline above: we would use a subset of R of size $1 + \log_p(m-1)$ as an assumed base, to avoid having to repeat the procedure with different guesses as the base, and use some sort of hashing technique to check whether we had seen each new base image before; however, although hashing has good average case complexity, the worst-case asymptotic complexity of this method would not be nearly linear, and the aim in this section is to obtain a nearly linear algorithm.

For each block system found by the above procedure we can then find the action on cosets of the kernel of the action of G_ω on the system, using the same methods as in Section 8. The procedure given earlier in this section, which assumed knowledge of the subgroup $G_{\alpha,\beta}$, can then be used. The implementation issues for this part of the algorithm are essentially the same as those for the equivalent part of the algorithm for affine groups. In Section 7 we obtained a deterministic nearly linear-time algorithm for this problem, and a similar result is easily obtained here in the same manner. Of course, the remarks made there about finding shallow Schreier trees and bounding $\#G$ in terms of the degree n to give an algorithm that is linear up to a factor that is a poly-logarithmic function of n alone (instead of n and $\#G$, as is usual with these 'nearly linear' algorithms) apply here as well.

We therefore have a deterministic nearly linear algorithm for recognising groups of class $\mathcal{L}'$ if we know the orbits of G_ω in advance, and thus we get a Monte-Carlo nearly linear implementation with no prerequisites. Using the Schreier generators as generators for G_ω yields a deterministic $O(sn^2)$ algorithm.

The non-simple Ree group, $\mathrm{R}(3) \cong \mathrm{P\Gamma L}(2,8)$. We apply similar techniques as for the rest of the class $\mathcal{L}$ to the group $\mathrm{R}(3)$, in its 2-transitive representation on 28 points. Proposition 9.1 no longer holds, but there is a pair ξ such that the smallest block of G acting on $\Gamma^{(2)}$ containing ξ and (α, β) is the desired block. It can be shown (using GAP [**8**]) that when the action on ordered pairs has been found, every one of the four blocks of imprimitivity of size $m-1$ that arise in this way projects to a solution subset. Therefore, a suitable block can be found by repeated calls to Atkinson's algorithm, or similar.

As for the other Ree groups, $G = \mathrm{R}(3)$ has the property that G_ω acts regularly on a largest suborbit; the part of the procedure concerned with finding subgroups of index 2 in G_ω is the same as in the general case for $\mathcal{L}'$.

10. *Summary and practical results*

We have given a deterministic, $O(sn^2)$ solution to the problem of Specification 1.1, and a deterministic nearly linear algorithm which solves Specification 1.1 if the G_ω-orbits are known in advance.

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

Using the techniques of [1], this result leads to a Monte-Carlo algorithm, also in nearly linear time, for all except the alternating groups (in the natural action), since the collection of all other 2-homogeneous groups is a class of small-base groups. Note that this is a one-sided Monte-Carlo algorithm: if the algorithm finds an action on pairs then its output is always correct; it is only if it reports that no such action exists that there is a possibility of error.

The $O(sn^2)$ algorithm based on the full set of Schreier generators is spatially expensive ($\Omega(n^2)$), and so in practice it seems important to reduce the number of generators used for G_ω to a more manageable level, or to accept a slight decrease in asymptotic temporal efficiency to handle the Schreier generators in a spatially more efficient manner.

Experimental results. The procedures described in the foregoing sections have been implemented in GAP [8]. The GAP code for the implementation is made available in Appendix C. The important part of this paper is the deterministic algorithms that start with generators for G_ω, and it is only these that have been tested, using the GAP implementation of the Schreier–Sims procedure to provide generators for G_ω. The time to run the Schreier–Sims routine has not been included in these timings.

The program was tested on a 200MHz. Pentium with 32Mb. of main store, running GAP under the Linux operating system. Timings are given in seconds of C.P.U. time, as reported by GAP, averaged over several runs with different random generating sets, usually of size 2 or 3; each time, the generators were conjugated by a different randomly chosen permutation of the points of Ω, to hide the structure of Ω as a set of pairs.

As is to be expected, the routine for the class $\mathcal{Z}$, using adjacent-point sets, runs fastest. For the groups with small r_i, the times taken showed little variation for any particular group, and were approximately proportional to the input degree, with the groups $\mathrm{PSL}(8, 2)$ acting on pairs of input degree $\binom{255}{2}$, and Co_3 acting on pairs of input degree $\binom{276}{2}$, requiring 10 and 11 seconds respectively (in addition to the time spent in the Schreier–Sims procedure).

The running times for the other classes were significantly longer than these timings. Typical results include recognising the action on pairs of $\mathrm{ASL}(5, 3)$ (input degree $\binom{243}{2}$) in under 30 seconds, the action on pairs of $\mathrm{ASL}(2, 17)$ (input degree $\binom{289}{2}$) in under a minute and the action on pairs of $\mathrm{ASL}(3, 7)$ (input degree $\binom{343}{2}$) in 1.5 minutes (all these timings are the average over several runs of the algorithm). The even characteristic is slower, and requires more memory: $\mathrm{AGL}(1, 256)$ acting on pairs (input degree $\binom{256}{2}$) was recognised in about 2.5 minutes, and $\mathrm{ASL}(2, 16)$ (same input degree) in 4–5 minutes.

When the input is not an action on pairs, the program normally recognises this quickly, using the suborbit structure. However, it is possible to create examples that take some time. For example, (using ATLAS notation) the group $15 \times 2^8 : 17$ has order $65280 = 2\binom{256}{2}$. The action on the cosets of one of the subgroups of order 2 is therefore an action of degree $\binom{256}{2}$ with suborbits of length at most 2, and the program took nearly 2.5 minutes to establish that this was not an action on pairs.

Acknowledgements. I am indebted to my supervisor, Peter Neumann, for suggesting this problem and for assistance in preparation of this paper. I have also incorporated an improvement to my proof of Proposition 9.1 which he suggested, and also improvements to one of the proofs in Appendix B. I am grateful to the referees for many helpful comments and suggestions. My doctoral research is supported by the E.P.S.R.C.

Appendix A. *Orbits of* $\mathrm{Sp}(2d, 2)$

The groups $\mathrm{Sp}(2d, 2)$ each have two 2-transitive actions, and we study the suborbits of the actions on unordered pairs arising out of these actions. We prove that there are at most 9 suborbits and calculate their lengths. It follows immediately from the result of this appendix that $r_9(m) = 1$ (see Table 1); it is shown in Section 5 how this leads to an efficient test based on the TestSet subroutine.

We follow [**4**, p. 247ff.]. Let F be a field, $d \geqslant 1$ a fixed integer, and $V = F^{2d}$, the space of row vectors of length $2d$ over F. Define two block matrices over F as follows:

$$e = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \text{ and } f = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} = e - e^T$$

where 0 and 1 denote the $d \times d$ zero and identity matrices respectively. Then $\mathrm{Sp}(2m, F)$ is the group of all $2d \times 2d$ invertible matrices x such that $xfx^T = f$.

Associated with f is the antisymmetric bilinear form $\phi : V \times V \to F$ defined by $\phi(u, v) = ufv^T$. We consider quadratic forms $\theta : V \to F$ satisfying

$$\phi(u, v) = \theta(u + v) - \theta(u) - \theta(v) \tag{A.1}$$

for all $u, v \in V$. If char $F \neq 2$ then there is a unique such form, namely $\theta(u) = \frac{1}{2}\phi(u, u)$, but if char $F = 2$ then this relationship breaks down, and in general there may be many quadratic forms associated with a given bilinear form, or none at all.

In this case, let $F = \mathbb{F}_2$. We summarise results from [**4**]. It can be seen that the quadratic form

$$\theta_0(u) = ueu^T (= ue^T u^T)$$

satisfies (A.1) for the specific ϕ defined earlier. Let X denote the set of all quadratic forms satisfying this condition. Then it is shown that $X = \{\theta_a \mid a \in V\}$ where

$$\theta_a(u) = ueu^T + ufa^T = \theta_0(u) + \phi(u, a). \tag{A.2}$$

Now $G = \mathrm{Sp}(2d, F)$ acts on X by $\theta^g(u) = \theta(ug^{-1})$ where $\theta \in X$, $g \in G$ and $u \in V$. Define transvections $t_a \in G$ by $ut_a = u + \phi(u, a)a$. Then $t_a^{-1} = t_a$ and $g^{-1}t_a g = t_{ag}$ for all $g \in G$.

Lemma A.1. *[**4**, Lemma 7.7A]*

(i) *For all $a, c \in V$ we have*

$$\theta_a^{t_c} = \begin{cases} \theta_a & \text{if } \theta_a(c) = 1 \\ \theta_{a+c} & \text{if } \theta_a(c) = 0. \end{cases}$$

(ii) *For all $a, b \in V$ there is at most one $c \in V$ such that t_c maps θ_a to θ_b. Such a c exists if and only if $\theta_0(a) = \theta_0(b)$ (and then $c = a + b$).*

It follows that X splits into two G-orbits, namely

$$X^+ = \{\theta_a \mid \theta_0(a) = 0\}$$

and

$$X^- = \{\theta_a \mid \theta_0(a) = 1\}$$

of sizes $\#X^+ = 2^{d-1}(2^d + 1)$ and $\#X^- = 2^{d-1}(2^d - 1)$. It is then shown in [**4**] that G acts 2-transitively on X^+ and on X^-.

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

We use techniques similar to those in [4] to determine the orbits of the stabilizer of two points in the 2-transitive actions. We then do the same for the orbits of the various 3-point stabilizers in these actions, and deduce the suborbits in the action on unordered pairs.

In places our analysis will hold only for $d \geqslant 5$. It has to be checked by hand or GAP [8] that the same result is true for $2 \leqslant d \leqslant 4$.

The following lemma is a generalisation of Lemma 7.7B in [4], where it is part of the proof of 2-transitivity. For $a \in V$ and $\epsilon \in \mathbb{F}_2$ define $L(a, \epsilon) = \{v \in V \mid \phi(v, a) = \epsilon\}$. As in [4], if $a_1, \ldots, a_r \in V$ are linearly independent and $\epsilon_1, \ldots, \epsilon_r \in \mathbb{F}_2$ then $L(a_1, \epsilon_1) \cap \cdots \cap L(a_r, \epsilon_r) = U + w_0$ for some subspace U of V of dimension $2m - r$ and some $w_0 \in V$. When all $\epsilon_i = 0$ we have $w_0 = 0$.

Lemma A.2. *Let $a_1, \ldots, a_r$ be linearly independent elements of V, where $r < d$. Let $K = L(a_1, \epsilon_1) \cap \cdots \cap L(a_r, \epsilon_r)$ for some $\epsilon_1, \ldots, \epsilon_r \in \mathbb{F}_2$. Then the sets $K^+ = K \cap \{a \mid \theta_0(a) = 0\}$ and $K^- = K \cap \{a \mid \theta_0(a) = 1\}$ both have size at least 2^{2d-r-2}.*

Proof. Since the a_i are linearly independent, $U = L(a_1, 0) \cap \cdots \cap L(a_r, 0)$ is a subspace of V of dimension $2d - r > r$. It therefore contains an element b that is linearly independent of $a_1, \ldots, a_r$. Thus $K_0 = K \cap L(b, \epsilon)$ has size 2^{2d-r-1} for any $\epsilon \in \mathbb{F}_2$. Fix $\epsilon = \theta_0(b) + 1$. Let $w \in K_0$. By choice of b, both w and $w + b$ lie in K. In fact, they both lie in K_0, since $\phi(w + b, b) = \phi(w, b) + \phi(b, b)$ and $\phi(b, b) = 0$. On the other hand, $\theta_0(w + b) = \theta_0(w) + \theta_0(b) + \phi(w, b) = \theta_0(w) + \theta_0(b) + \epsilon = \theta_0(w) + 1$, so one of w and $w + b$ lies in K^+ and the other lies in K^-. This gives a pairing of the elements of K_0, since $(w + b) + b = w$, with one element of each pair in K^+ and the other in K^-. Therefore half of K_0 lies in K^+ and the other half in K^- and each of these sets contains at least 2^{2d-r-2} elements. □

Lemma A.3. *Suppose v_1, v_2, t, w are distinct elements of V, with $\theta_0(v_1) = \theta_0(v_2) = \theta_0(t) = \theta_0(w) = \epsilon$. If there is $g \in G$ fixing θ_{v_1} and θ_{v_2} and mapping θ_t to θ_w then $\phi(t, v_1 + v_2) = \phi(w, v_1 + v_2)$.*

There are 2^{2d-2} elements $t \in V$ with $\theta_0 b = \epsilon$ and $\phi(t, v_1 + v_2) = \theta_0(v_1 + v_2) + 1$, and $2^{d-1}(2^{d-1} + (-1)^\epsilon)$ elements $t \in V$ with $\theta_0 t = \epsilon$ and $\phi(t, v_1 + v_2) = \theta_0(v_1 + v_2)$.

Proof. Let $g \in G_{\theta_{v_1}, \theta_{v_2}}$. Then $\theta_{v_1}^g(u) + \theta_{v_2}^g(u) = \theta_{v_1}(u) + \theta_{v_2}(u)$ for all $u \in V$ and so by (A.2) and the definition of the action of G,

$$\phi(ug^{-1}, v_1 + v_2) = \phi(u, v_1 + v_2) \text{ for all } u \in V. \tag{A.3}$$

Now suppose $\theta_t^g = \theta_w$. For $u \in V$ we have $\theta_w(u) = \theta_{v_1 + (v_1 + t)}(ug^{-1}) = \theta_{v_1}(ug^{-1}) + \phi(ug^{-1}, a_1 + t)$ by two applications of (A.2). Then $\theta_{v_1}(ug^{-1}) = \theta_{v_1}(u)$ and $\phi(ug^{-1}, v_1 + t) = \phi(u, (v_1 + t)g)$ since $g \in \mathrm{Sp}(2d, 2)$. Therefore (by further applications of (A.2)) $\theta_w = \theta_{v_1 + (v_1 + t)g}$, and so $w = v_1 + (v_1 + t)g$. Now $\phi(v_1 + (v_1 + t)g, v_1 + v_2) = \phi(v_1, v_1 + v_2) + \phi((v_1 + t)g, v_1 + v_2)$ and by (A.3) applied to the second term this equals $\phi(t, v_1 + v_2)$.

For the second part, consider Lemma A.2 with $r = 0$. In the proof, we get $U = V$, and we can certainly take b to be $v_1 + v_2$ as v_1, v_2 are distinct. Then $K_0 = L(v_1 + v_2, \theta_0(v_1 + v_2) + 1)$ contains 2^{2d-2} vectors v with $\theta_0(v) = 0$ and 2^{2d-2} with $\theta_0(v) = 1$, so the set of points $t \in V$ with $\theta_0(t) = \epsilon$ and $\phi(t, v_1 + v_2) = \theta_0(v_1 + v_2) + 1$ has size 2^{2d-2}. Since there are $2^{d-1}(2^d + (-1)^\epsilon)$ elements $t \in V$ with $\theta_0(t) = \epsilon$, it follows that there are $2^{d-1}(2^d + (-1)^\epsilon) - 2^{2d-2} = 2^{d-1}(2^{d-1} + (-1)^\epsilon)$ points $v \in V$ with $\theta_0(v) = \epsilon$ and $\phi(t, v_1 + v_2) = \theta_0(v_1 + v_2)$. □

Proposition A.4. *Let $d \geqslant 4$. Suppose a_1, a_2, b, c are distinct elements of V, with $\theta_0(a_1) = \theta_0(a_2) = \theta_0(b) = \theta_0(c) = \epsilon$. Then if $\phi(b, a_1 + a_2) = \phi(c, a_1 + a_2)$ there is $g \in G$ fixing θ_{a_1} and θ_{a_2} and mapping θ_b to θ_c.*

Proof. We show that there exists $w \in V$ such that $\theta_0(w) = \epsilon$ and

$$\theta_{a_1}(b + w) = \theta_{a_1}(c + w) = \theta_{a_2}(b + w) = \theta_{a_2}(c + w) = 1. \tag{A.4}$$

It will then follow from Lemma A.1 that $g = t_{b+w}t_{c+w}$ has the desired properties.

We have $\theta_{a_1}(b + w) = \theta_{a_1}(b) + \theta_{a_1}(w) + \phi(w, b) = \theta_0(b) + \theta_0(w) + \phi(b, a_1) + \phi(w, a_1) + \phi(w, b)$ and similarly for the others, so the conditions (A.4) are equivalent to

$$\phi(w, a_1 + b) = 1 + \phi(b, a_1),\ \ \phi(w, a_1 + c) = 1 + \phi(c, a_1),$$
$$\phi(w, a_2 + b) = 1 + \phi(b, a_2),\ \ \phi(w, a_2 + c) = 1 + \phi(c, a_2).$$

The set of $w \in V$ satisfying these conditions is

$$K = L(a_1 + b, 1 + \phi(b, a_1)) \cap L(a_1 + c, 1 + \phi(c, a_1)) \cap$$
$$L(a_2 + b, 1 + \phi(b, a_2)) \cap L(a_2 + c, 1 + \phi(c, a_2)).$$

Now $a_2 + c = (a_1 + b) + (a_1 + c) + (a_2 + b)$ and since $\phi(b, a_1 + a_2) = \phi(c, a_1 + a_2)$, we have $1 + \phi(c, a_2) = (1 + \phi(b, a_1)) + (1 + \phi(c, a_1)) + (1 + \phi(b, a_2))$. Therefore

$$K = L(a_1 + b, 1 + \phi(b, a_1)) \cap L(a_1 + c, 1 + \phi(c, a_1)) \cap L(a_2 + b, 1 + \phi(b, a_2)).$$

If it can be shown that $a_1 + b$, $a_1 + c$ and $a_2 + b$ are linearly independent, then it will follow by Lemma A.2 that θ_0 is non-constant on K, and the proof will be complete.

Clearly $a_1 + b$ and $a_2 + b$ are linearly independent since a_1, a_2 and b are distinct (and $\mathbb{F}_2$ has only 2 elements). Similarly, $a_1 + c$ is non-zero, and not equal to $a_1 + b$ or $a_1 + a_2 = ((a_1 + b) + (a_2 + b))$. Therefore either $a_1 + b$, $a_1 + c$ and $a_2 + b$ are linearly independent, or $a_1 + c = a_2 + b$. So we certainly have a proof except in the case where $c = a_1 + a_2 + b$. In this case there is always some $x \in V$ with $\theta_0(x) = \epsilon$, $\phi(x, a_1 + a_2) = \phi(b, a_1 + a_2) = \phi(c, a_1 + a_2)$ and $x \notin \{a_1, a_2, b, c\}$, since (by the preceding lemma) there are at least $2^{2d-2} - 2^{d-1}$ elements fulfilling the first two of these conditions, and $d \geqslant 4$ by hypothesis so this number is larger than 4. Now by the proof in the linearly independent case there exist $g_1, g_2 \in G$ that both fix θ_{a_1} and θ_{a_2}, and such that g_1 maps θ_b to θ_x, whilst g_2 maps θ_x to $\theta_c = \theta_{a_1+a_2+b}$. Then $g = g_1g_2$ will suffice. □

Corollary A.5. *In the action of G on $\Gamma = X^{\pm}$, the two-point stabilizer $G_{\alpha,\beta}$ has four orbits: $\{\alpha\}$, $\{\beta\}$, one of size $2^{2(d-1)}$ and one of size $2(2^{d-2} + (-1)^{\epsilon})(2^{d-1} - (-1)^{\epsilon})$, where $\epsilon = 0$ if $\Gamma = X^+$ and $\epsilon = 1$ if $\Gamma = X^-$.*

Proof. Let $\alpha = \theta_{a_1}$ and $\beta = \theta_{a_2}$ where $\theta_0(a_1) = \theta_0(a_2) = \epsilon$. Note that $\phi(a_i, a_1 + a_2) = \theta_0(a_1 + a_2)$ for $i = 1, 2$. The result now follows by the preceding proposition and lemma, since $2^{d-1}(2^{d-1} + (-1)^{\epsilon}) - 2 = 2(2^{d-2} + (-1)^{\epsilon})(2^{d-1} - (-1)^{\epsilon})$. □

We can calculate the orbits of any 3-point stabilizer, as well.

Proposition A.6. *Let $a_1, a_2, a_3 \in V$, with $\theta_0(a_i) = \epsilon$ for $i = 1, 2, 3$. Let $\epsilon_0 = \phi(a_3, a_1 + a_2) + \theta_0(a_1 + a_2)$. Then the subset $\{v \in V \mid \theta_0(v) = \epsilon\}$ of V is partitioned by the sets $X_{\lambda\mu}$, where $\lambda, \mu \in \mathbb{F}_2$ and*

$$X_{\lambda\mu} = \{v \in V \mid \theta_0(v) = \epsilon\} \cap L(a_1 + a_2, \theta_0(a_1 + a_2) + \lambda + \mu + \epsilon_0)$$

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

$$\cap\, L(a_1 + a_3, \theta_0(a_1 + a_3) + \lambda)$$
$$\cap\, L(a_2 + a_3, \theta_0(a_2 + a_3) + \mu).$$

If $\epsilon_0 = 0$ then $\#X_{00} = 2^{d-1}(2^{d-2} + (-1)^\epsilon)$ and $\#X_{01} = \#X_{10} = \#X_{11} = 2^{2d-3}$. If $\epsilon_0 = 1$ then $\#X_{00} = \#X_{01} = \#X_{10} = 2^{d-2}(2^{d-1} + (-1)^\epsilon)$ and $\#X_{11} = 2^{d-2}(2^{d-1} - (-1)^\epsilon)$.

Proof. Note that $\phi(u, a_1 + a_2) = \phi(u, a_1 + a_3) + \phi(u, a_2 + a_3)$ for $u \in V$, and that $\theta_0(a_1 + a_2) = \theta_0(a_1 + a_3) + \theta_0(a_2 + a_3) + \phi(a_1 + a_3, a_2 + a_3) = \theta_0(a_1 + a_3) + \theta_0(a_2 + a_3) + \epsilon_0$. It follows that

$$\begin{aligned} X_{\lambda\mu} = \{v \in V \mid \theta_0(v) = \epsilon\} &\cap L(a_1 + a_3, \theta_0(a_1 + a_3) + \lambda) \\ &\cap L(a_2 + a_3, \theta_0(a_2 + a_3) + \mu) \end{aligned}$$

and so the $X_{\lambda\mu}$ do indeed partition $\{v \in V \mid \theta_0(v) = \epsilon\}$.

Suppose $\epsilon_0 = 0$. Then

$$\begin{aligned} X_{00} \cup X_{01} &= \{v \in V \mid \theta_0(v) = \epsilon\} \cap L(a_1 + a_3, \theta_0(a_1 + a_3)) \\ X_{00} \cup X_{10} &= \{v \in V \mid \theta_0(v) = \epsilon\} \cap L(a_2 + a_3, \theta_0(a_2 + a_3)) \\ X_{00} \cup X_{11} &= \{v \in V \mid \theta_0(v) = \epsilon\} \cap L(a_1 + a_2, \theta_0(a_1 + a_2)) \end{aligned}$$

and each of these sets has size $2^{d-1}(2^{d-1} + (-1)^\epsilon)$ by Lemma A.3. Also

$$X_{01} \cup X_{10} = \{v \in V \mid \theta_0(v) = \epsilon\} \cap L(a_1 + a_2, \theta_0(a_1 + a_2) + 1)$$

and similar expressions for $X_{01} \cup X_{11}$ and $X_{10} \cup X_{11}$. Each of these sets has size 2^{2d-2} by Lemma A.3. It follows that $\#X_{01} = \#X_{10} = \#X_{11} = 2^{2d-3}$ and $\#X_{00} = 2^{d-1}(2^{d-1} + (-1)^\epsilon - \#X_{01} = 2^{d-1}(2^{d-2} + (-1)^\epsilon)$. The $\epsilon_0 = 1$ case is entirely analogous. □

Proposition A.7. *Let $d \geqslant 5$. Let $a_1, a_2, a_3 \in V$, with $\theta_0(a_i) = \epsilon$ for $i = 1, 2, 3$. Let ϵ_0 and the sets $X_{\lambda\mu}$ be as in the previous proposition. Then the sets*

$$\Xi_{\lambda\mu} = \{\theta_x \mid x \in X_{\lambda\mu} \setminus \{a_1, a_2, a_3, a_1 + a_2 + a_3\}\}$$

form $G_{\theta_{a_1}, \theta_{a_2}, \theta_{a_3}}$-orbits.

Proof. First note that $\theta_{a_1+a_2+a_3}(u) = \theta_{a_1}(u) + \phi(u, a_2 + a_3)$ for all $u \in V$ and so by (A.3), $\theta_{a_1+a_2+a_3}$ must be fixed by $G_{\theta_{a_1}, \theta_{a_2}, \theta_{a_3}}$. A consequence of Lemma A.3 is therefore that the sets $\Xi_{\lambda\mu}$ must be unions of $G_{\theta_{a_1}, \theta_{a_2}, \theta_{a_3}}$-orbits.

Fix $\lambda, \mu \in \mathbb{F}_2$, and let $b, c \in X_{\lambda\mu}$, with $b, c \notin \{a_1, a_2, a_3, a_1 + a_2 + a_3\}$. As in Proposition A.4, we show that there exists $w \in V$ such that $\theta_0(w) = \epsilon$ and

$$\theta_{a_i}(b + w) = \theta_{a_i}(c + w) = 1 \tag{A.5}$$

for $i = 1, 2, 3$. It will then follow from Lemma A.1 that $g = t_{b+w} t_{c+w}$ has the desired properties.

As in Proposition A.4, the equations (A.5) are equivalent to

$$\phi(w, a_i + b) = 1 + \phi(b, a_i), \ \ \phi(w, a_i + c) = 1 + \phi(c, a_i)$$

for $i = 1, 2, 3$. The set of elements of V satisfying these conditions is

$$\begin{aligned} K = L(a_1 + b, 1 + \phi(b, a_1)) &\cap L(a_2 + b, 1 + \phi(b, a_2)) \\ \cap\, L(a_3 + b, 1 + \phi(b, a_3)) &\cap L(a_1 + c, 1 + \phi(c, a_1)) \end{aligned}$$

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

since $a_j + c = (a_1 + b) + (a_j + b) + (a_1 + c)$ and $\phi(c, a_j) = \phi(b, a_1) + \phi(b, a_j) + \phi(c, a_1)$ for $j = 2, 3$, because b and c are in the same set $X_{\lambda\mu}$.

In the case where $a_1 + b$, $a_2 + b$, $a_3 + b$ and $a_1 + c$ are linearly independent, the result now follows by Lemma A.2, since $d \geqslant 5$. Now $a_1 + b$, $a_2 + b$ and $a_3 + b$ must be linearly independent, since a_1, a_2, a_3 and b are distinct and $b \neq a_1 + a_2 + a_3$. Similarly $a_1 + c$, $a_2 + c$ and $a_3 + c$ are linearly independent. If $a_1 + b$, $a_2 + b$, $a_3 + b$ and $a_1 + c$ are not linearly independent, they span a space of dimension at most 3, and so (since $X_{\lambda\mu}$ has size at least $2^{2d-3} - 2^{d-1} > 8$) there is an element $x \in X_{\lambda\mu}$ such that $a_1 + x$ is linearly independent of $a_1 + b$, $a_2 + b$, $a_3 + b$ and $a_1 + c$, and therefore both $\{a_1 + b, a_2 + b, a_3 + b, a_1 + x\}$ and $\{a_1 + c, a_2 + c, a_3 + c, a_1 + x\}$ are linearly independent. We have therefore shown that we can map from θ_b to θ_x, and from θ_c to θ_x within $G_{\theta_{a_1}, \theta_{a_2}, \theta_{a_3}}$ and so the proof is complete. □

Corollary A.8. *If* $\epsilon_0 = 0$ *then there are eight* $G_{\theta_{a_1}, \theta_{a_2}, \theta_{a_3}}$*-orbits: four of size* 1*, three of size* 2^{2d-3} *and one of size* $2^{d-1}(2^{d-2} + (-1)^{\epsilon}) - 4 = 4(2^{d-2} - (-1)^{\epsilon})(2^{d-3} + (-1)^{\epsilon})$*. If* $\epsilon_0 = 1$ *then there are seven* $G_{\theta_{a_1}, \theta_{a_2}, \theta_{a_3}}$*-orbits: three of size* 1*, three of size* $2^{d-2}(2^{d-1} + (-1)^{\epsilon}) - 1 = (2^{d-2} + (-1)^{\epsilon})(2^{d-1} - (-1)^{\epsilon})$ *and one of size* $2^{d-2}(2^{d-1} - (-1)^{\epsilon})$.

Proof. Firstly $\theta_0(a_1 + a_2 + a_3) = \theta_0(a_3) + \theta_0(a_1 + a_2) + \phi(a_3, a_1 + a_2) = \epsilon + \epsilon_0$. So $\theta_{a_1+a_2+a_3}$ lies in the same G-orbit as θ_{a_i} if and only if $\epsilon_0 = 0$. Taken with the three orbits $\{a_i\}$, this gives the singleton orbits listed.

Secondly, $\phi(a_i, a_i + a_j) = \theta_0(a_i + a_j)$ for distinct $i, j \in \{1, 2, 3\}$. This means that if $\epsilon_0 = 0$ then $a_i \in X_{00}$ for $i = 1, 2, 3$, and if $\epsilon_0 = 1$ then $a_1 \in X_{01}$, $a_2 \in X_{10}$ and $a_3 \in X_{00}$. We also have $\phi(a_1 + a_2 + a_3, a_i + a_j) = \phi(a_k, a_i + a_j)$ where $\{i, j, k\} = \{1, 2, 3\}$ and $\phi(a_k, a_i + a_j) = \theta_0(a_i + a_j) + \epsilon_0$ (e.g., $\phi(a_1, a_2 + a_3) = \theta_0(a_2 + a_3) + (\phi(a_3, a_1 + a_2) + \theta_0(a_1 + a_2))$). Therefore if $\epsilon_0 = 0$, $a_1 + a_2 + a_3 \in X_{00}$. The results follow from the previous proposition and the sizes of the sets $X_{\lambda\mu}$ as calculated in Proposition A.6. □

We can now calculate the sizes of the suborbits of G acting on unordered pairs. Let $\alpha = \theta_{a_1}$, $\beta = \theta_{a_2}$ where $\theta_0(a_i) = \epsilon$ $(i = 1, 2)$. Let $\omega = \{\alpha, \beta\}$, and consider the action of G_ω on $\Omega = \{\theta_v \mid \theta_0(v) = \epsilon\}^{\{2\}}$. Observe that G_ω is generated by $G_{\theta_{a_1}, \theta_{a_2}}$ and the transvection $h = t_{a_1+a_2}$. Let $B = \{b \mid \theta_0(b) = \epsilon,\ \phi(b, a_1 + a_2) + \theta_0(a_1 + a_2) = 0,\ b \notin \{a_1, a_2\}\}$, and $C = \{c \mid \theta_0(c) = \epsilon,\ \phi(c, a_1 + a_2) + \theta_0(a_1 + a_2) = 1,\ c \notin \{a_1, a_2\}\}$. By Lemma A.1, $\theta_b^h = \theta_{b+a_1+a_2}$ for $b \in B$, and $\theta_c^h = \theta_c$ for $c \in C$.

It is easy to see that the sets $\{\omega\}$, $\{\{\theta_{a_i}, \theta_b\} \mid i \in \{1, 2\},\ b \in B\}$ and $\{\{\theta_{a_i}, \theta_c\} \mid i \in \{1, 2\},\ c \in C\}$ are G_ω-orbits, of sizes 1, $2\#B$ and $2\#C$ respectively, and that the adjacent-point set $\Lambda_\Gamma(\omega)$ is the union of these three sets.

For the remaining orbits we look first at orbits on *ordered* pairs. By Proposition A.7 and its corollary, there are three $G_{\theta_{a_1}, \theta_{a_2}}$-orbits of ordered pairs $(\theta_{b_1}, \theta_{b_2})$ where $b_1, b_2 \in B$, of sizes $\#B$, $\#B.\#\Xi_{00}$ and $\#B.\#\Xi_{11}$, since if $a_3 \in B$ then $\epsilon_0 = 0$ and $B \subseteq X_{00} \cup X_{11}$. The orbit containing $(\theta_{b_2}, \theta_{b_1})$ must have the same size as that containing $(\theta_{b_1}, \theta_{b_2})$, and so, since these three orbits have different sizes, each must be symmetric, that is, contain the pair $(\theta_{b_2}, \theta_{b_1})$ for every pair $(\theta_{b_1}, \theta_{b_2})$ lying in it. Therefore these three orbits correspond to three orbits of unordered pairs, each half the size. Because $b + a_1 + a_2 \in B$ whenever b does, if any of these orbits were not also G_ω-orbits, then there would have to be a pair of them, such that the union of the pair was a G_ω-orbit. The two orbits would have to have the same size for this to happen, and this is not the case, so each is itself a G_ω-orbit.

There are two $G_{\theta_{a_1},\theta_{a_2}}$-orbits of ordered pairs $(\theta_{c_1}, \theta_{c_2})$ where $c_1, c_2 \in C$, of sizes $\#C.\#\Xi_{00}$ and $\#C.\#\Xi_{11}$. As before, these have different sizes so they are symmetric, and there are two corresponding orbits on unordered pairs each of half the size of its ‘image’. They are also G_ω-orbits, since h fixes θ_c for $c \in C$.

Table A.1: Suborbit Lengths of $\mathrm{Sp}(2d, 2)$ acting on pairs

1
$2^2(2^{d-2} + (-1)^\epsilon)(2^{d-1} - (-1)^\epsilon)$
2^{2d-1}
$(2^{d-2} + (-1)^\epsilon)(2^{d-1} - (-1)^\epsilon)$
$2^{2d-3}(2^{d-2} + (-1)^\epsilon)(2^{d-1} - (-1)^\epsilon)$
$4(2^{d-2} - (-1)^\epsilon)(2^{d-3} + (-1)^\epsilon)(2^{d-2} + (-1)^\epsilon)(2^{d-1} - (-1)^\epsilon)$
$2^{2d-3}(2^{d-2} + (-1)^\epsilon)(2^{d-1} - (-1)^\epsilon)$
$2^{2d-3}2^{d-2}(2^{d-1} - (-1)^\epsilon)$
$2^{2d-1}(2^{d-2} + (-1)^\epsilon)(2^{d-1} - (-1)^\epsilon)$

Finally, there are two $G_{\theta_{a_1},\theta_{a_2}}$-orbits of ordered pairs (θ_c, θ_b) where $c \in C$, $b \in B$, of sizes $\#C.\#\Xi_{01}$ and $\#C.\#\Xi_{10}$. These are clearly not symmetric, and the corresponding orbits of unordered pairs have the same sizes as their ‘images’. Note that the orbits are of the same size as each other, and that if $a_3 = c$ and $b \in X_{01}$ then $b + a_1 + a_2 \in X_{10}$ since $\epsilon_0 = 1$. Therefore the union of the two orbits is a G_ω-orbit.

A summary table of the orbit lengths of G_ω acting on Ω is included in Table A.1. Checks using GAP have shown that the table is correct in the cases $2 \leqslant d \leqslant 4$, although in the $d = 2$ cases and one of the $d = 3$ cases there are fewer suborbits, as one or more of the orbit lengths given in the table evaluate to 0.

Appendix B. *Bases for Groups of Lie Rank* 1

This appendix contains results about the groups of class $\mathcal{L}'$, that is, the groups of Lie rank 1. If (G, Γ) is the 2-transitive action of such a group then we consider the action of G on $\Omega = \Gamma^{\{2\}}$; more precisely, we take $\omega \in \Omega$ and look at the action of Ω on a largest G_ω-orbit R. We show that this action is faithful and either regular or very nearly regular, and we examine how to extend a set containing ω and one element of R to form a base for G. (Recall that a base for a permutation group is a subset of the domain whose point-wise stabilizer is trivial.)

These results are used in Section 9 to provide a means of finding the subgroups of G_ω of index 2. Recall from that section that $\mathcal{L}$ denotes the 2-transitive groups whose socle is one of $\mathrm{PSU}(3, q)$ (for q a prime power greater than 2), $\mathrm{Sz}(q)$ (for q an odd power of 2 greater than 2) and $\mathrm{R}(q)$ (for q an odd power of 3 greater than 3). The class $\mathcal{L}'$ contains these groups and those with socle $\mathrm{PSL}(2, q)$ for q a prime power greater than 3.

Proposition B.1. *Let* (G, Γ) *lie in the class* $\mathcal{L}$*, and let* $\Omega = \Gamma^{\{2\}}$*. Then there exists a base for* (G, Ω) *of size* 2*.*

Proof. As in Section 9, we refer the reader to [4] for a description of these groups.

Let K denote the field of q elements (q^2 for the $\mathrm{PSU}(3,q)$ case, when we denote the fixed field of the automorphism $x \mapsto \bar{x}$ of order 2 of K by K_0), and let p be the prime dividing q. Let S denote the socle of the group G.

As was mentioned in Section 9, the stabilizer S_α is the semidirect product of a Sylow p-subgroup T of S by the two-point stabilizer $S_{\alpha,\beta}$, which we denote by H. The elements of T can be parametrised by 2 or 3 parameters from K: for $\mathrm{PSU}(3,q)$, $T = \{t_{a,b} \bigm| a, b \in K,\ a\bar{a} + b + \bar{b} = 0\}$ (note that the parameters here are interchanged with respect to those in [4] to be consistent with the other groups); for $\mathrm{Sz}(q)$, $T = \{t_{a,b} \bigm| a, b \in K\}$ (of course, the $t_{a,b}$ have completely different interpretations in the two cases) and for $\mathrm{R}(q)$, $T = \{t_{a,b,c} \bigm| a, b, c \in K\}$. Note that for any automorphism σ^* of one of these groups that is induced by a field automorphism σ of K, we have $t_{a,b}^{\sigma^*} = t_{a^\sigma,b^\sigma}$ (and similarly for $t_{a,b,c}$).

The cyclic group H has in all cases a natural identification with either the multiplicative group $K^\times$ of K, or the subgroup of index 3 in $K^\times$; the latter situation arises if and only if S is $\mathrm{PSU}(3,q)$ and $3 \mid q+1$. Thus we can write $H = \{\eta_\kappa \mid \kappa \in H_0\}$ for some subgroup H_0 of $K^\times$. The action by conjugation of H on T satisfies $t_{a,b}^{\eta_\kappa} = t_{a\kappa,b'}$ (or $t_{a,b,c}^{\eta_\kappa} = t_{a\kappa,b',c'}$ as applicable) for some $b', c' \in K$.

It is sufficient to find points $\alpha, \beta, \gamma \in \Gamma$ such that $G_{\alpha,\beta,\gamma} = 1$ since then $\{\{\alpha,\beta\},\{\alpha,\gamma\}\}$ is a base for (G,Ω). Because (G,Γ) is 2-transitive, we can take α, β to be arbitrary points of Γ, as before, and use the notation introduced earlier. The problem then reduces to that of finding an element of T that is moved by every non-trivial element of $G_{\alpha,\beta}$ in the action on T by conjugation.

If we can find an element b of K that is moved by all automorphisms of K, and (if S is $\mathrm{PSU}(3,q)$) for which $t_{1,b}$ (or $t_{1,b,c}$ for some c) is an element of T, then we have done. This is because for any $h \in H$ or, in the $\mathrm{PSU}(3,q)$ case where $3 \mid q+1$, any product h of a diagonal automorphism and an element of H, the conjugate $t_{1,b}^h$ (or $t_{1,b,c}^h$) is not of the form $t_{1,b'}$ (or $t_{1,b',c'}$) for any b' (or b', c') in K. Since every element of $G_{\alpha,\beta}$ may be written as a product of a field automorphism followed by such an element h, this means that the element $t_{1,b}$ (or $t_{1,b,c}$) is moved by every non-trivial element of $G_{\alpha,\beta}$.

There are always elements of K moved by every field automorphism: just take any element b that generates K over the prime field. If it is fixed by a field automorphism, then it must lie in a proper subfield of K, contradicting the choice of b. Therefore we are done when S is a Suzuki or Ree group.

For the unitary case, we require an element $b \in K$ that is moved by every field automorphism, and for which $1 + b + \bar{b} = 0$. This is equivalent to showing that there is an element b in K that does not lie in any proper subfield of K and for which $1 + b + \bar{b} = 0$. The number of elements b of K with $1 + b + \bar{b} = 0$ is q, since the trace map $K \to K_0$ given by $x \mapsto x + \bar{x}$ is surjective and K_0-linear. Let N be the number of these elements that lie in proper subfields of K. We will show that $N < q$.

Let r be the integer such that $q = p^r$. Then K has size p^{2r} and the maximal subfields of K are K_0 and the fields of size $p^{2r/s}$ for odd prime divisors s of r. If $b \in K_0$ then $\bar{b} = b$; if in addition $1 + b + \bar{b} = 0$ then $1 + 2b = 0$ so this is impossible if $p = 2$ and otherwise there is precisely one such b. The number of b with $1 + b + \bar{b} = 0$ that lie in the field of size $p^{2r/s}$ is at most $p^{2r/s} - 1$, since 0 is not a solution to this equation. Therefore

$$N \leqslant 1 + \sum (p^{2r/s} - 1)$$

where the sum is over all odd prime divisors s of r. Therefore if r is a power of 2 we have $N \leqslant 1$, otherwise $N \leqslant tp^{2r/3}$ where t is the number of odd prime divisors of r. It is easy to see that $t < p^{r/3}$, so $N < p^r = q$. This concludes the proof in the unitary case. □

The remainder of this appendix is concerned with groups whose socle is $\mathrm{PSL}(2, q)$.

Lemma B.2. *Let $q = p^r$ where p is prime, and let Σ be a group of field automorphisms of $\mathbb{F}_q$, acting naturally on $\mathbb{F}_q$. Then there exists $a \in \mathbb{F}_q$ such that $\Sigma_a = 1$. Furthermore if $q \notin \{2, 4\}$ then a may be chosen so that a^{-1} does not lie in the same Σ-orbit as a. If q is odd then a is not a square in $\mathbb{F}_q^\times$, but in this case if $q \notin \{3, 5, 9\}$ then there also exists a square $b \in \mathbb{F}_q$ with $\Sigma_b = 1$ and $b^{-1} \notin b^\Sigma$.*

Proof. Let a generate $\mathbb{F}_q^\times$. Then a has order $p^r - 1$ and $a^{-1} = a^{p^r-2}$. Let $\sigma \in \Sigma \setminus \{1\}$. Then $\sigma(a) = a^{p^e}$ where $1 < e < r$. Clearly none of $p, p^2, \ldots, p^{r-1}$ is congruent to 1 mod $p^r - 1$, so $\Sigma_a = 1$. Also -1 is only congruent mod $p^r - 1$ to a member of the set $\{1, p, p^2, \ldots, p^{r-1}\}$ if $p = 2$ and r is 1 or 2, so a^{-1} does not lie in the orbit a^Σ unless q is 2 or 4.

Now assume q is odd, so a is a non-square satisfying both conditions. Let $b = a^2$. Then b has order $(p^r - 1)/2$ and if $p^e \equiv 1 \bmod (p^r - 1)/2$ for some e with $1 < e < r$ then $(p^r - 1)/2 + 1 = p^e$ (since $(p^r - 1) + 1$ is too big). Then $2p^e = p^r + 1$. Since $p^r \geqslant 3p^e$, this is impossible. Finally, if $p^e \equiv -1 \bmod (p^r - 1)/2$ for some e with $0 \leqslant e < r$ then $p^e = (p^r - 3)/2$ or $p^e = p^r - 2$. In the first case $2p^e + 3 = p^r$ and as $p^r \geqslant 3p^e$ this requires $p^e \leqslant 3$, and so $p^r \leqslant 9$. Therefore $r \leqslant 2$ and so either $e = 0$ and $p^r = 5$, or $e = 1$, $r = 2$ and $p^r = 9$. In the second case, as p is odd we must have $p^r = 3$. Therefore if $q \notin \{3, 5, 9\}$ then b is a square satisfying both conditions. □

Proposition B.3. *Let (G, Γ) be a group with socle $\mathrm{PSL}(2, q)$ (where $q > 3$) acting on projective points (so $\#\Gamma = m = q + 1$), and let $\Omega = \Gamma^{\{2\}}$. If (G, Γ) is not 3-transitive then there exists a base for (G, Ω) of size 2.*

Proof. We can regard Γ as the projective line $\mathbb{F}_q \cup \{\infty\}$, and the elements of $\mathrm{P\Gamma L}$ as transformations $x \mapsto (ax + b)/(cx + d)$ composed with field automorphisms acting naturally.

Suppose G is not 3-transitive on Γ. Then p is odd and $G_{0,\infty}$ has two orbits on $\Gamma \setminus \{0, \infty\}$, namely the set of squares in $\mathbb{F}_q^\times$ and the set of non-squares. Field automorphisms map squares to squares, but the non-trivial diagonal automorphism maps squares to non-squares and vice-versa. Thus if the group contains any outer automorphism not contained in the group of field automorphisms then it is 3-transitive on Γ. So we may reduce the problem to the case $\mathrm{P\Sigma L}(2, q) = \mathrm{PSL}(2, q) \rtimes \langle\phi\rangle$ where ϕ is the automorphism induced by the Frobenius map $x \mapsto x^p$ on F, and q is odd, since any non-3-transitive group of $\mathrm{PSL}(2, q)$-type is contained in this group, and so it suffices to find a base of size 2 for this group acting on $\Omega = \Gamma^{\{2\}}$.

We take $\alpha = \infty$ and $\beta = 0$. Our strategy will be to fix $\omega = \{\alpha, \beta\}$ and choose a suitable point ω_1, then show that $\{\omega, \omega_1\}$ is a base by taking an arbitrary element $h \in G$ fixing ω and showing that if h fixes ω_1 as well then $h = 1$. So let $h \in G_\omega$. Then there exist $t \in \mathbb{F}_q$ and $i \geqslant 0$ such that $h = g\phi^i$ or $h = gj\phi^i$ where $g : x \mapsto t^2 x$ and $j : x \mapsto -x^{-1}$.

We distinguish two cases, depending on whether $q \equiv 1 \bmod 4$ or $q \equiv 3 \bmod 4$. In the first case, -1 is a square in $\mathbb{F}_q$. By Lemma B.2, there exists $a \in \mathbb{F}_q^\times$ which is a non-square fixed only by the trivial element of $\langle\phi\rangle$, and such that a^{-1} does not occur in the orbit of a under $\langle\phi\rangle$. Let $\omega_1 = \{1, a\}$. Then $\omega_1^{\phi^i} = \{1, a_1\}$ where $a_1 \neq a, a^{-1}$ unless $\phi^i = 1$, and also a_1 is not a square. Then $\omega_1^{g\phi^i} = \{t^2, t^2 a_1\}$ and is only equal to ω_1 if $t^2 = 1$ and $t^2 a_1 = a$, since t^2 and 1 are squares and t^2a_1 and a are not. Thus $\omega_1^{g\phi^i} = \omega_1$ only in the case where $g = 1$ and $\phi^i = 1$. Also, $\omega_1^{gj\phi^i} = \{-t^2, -t^2a_1^{-1}\}$ and is only equal to ω_1 if $-t^2 = 1$ and $-t^2a_1^{-1} = a$, since 1 and $-t^2$ are squares and the other two elements are not. Thus $\omega_1^{gj\phi^i} = \omega_1$ only when $a = a_1^{-1}$, which is impossible. Thus $\{\omega, \omega_1\}$ is a base.

In the second case, -1 is not a square in $\mathbb{F}_q$. By Lemma B.2 there is a square $b \in F$ that is not fixed by any non-trivial element of $\langle\phi\rangle$ and is such that $b^{-1} \notin b^{\langle\phi\rangle}$. Let $\omega_1 = \{1, b\}$. Then $\omega_1^{\phi^i} = \{1, b_1\}$ where $b^{-1} \neq b_1$ and $b = b_1$ only if $\phi^i = 1$. Then $\omega_1^{g\phi^i} = \{t^2, t^2b_1\}$, which equals ω_1 only if the ratio of the two elements is the same, that is to say only if $b_1 \in \{b, b^{-1}\}$. Thus to fix ω_1 we would need $b_1 = b$ and so $\phi^i = 1$, and then we would also need $t^2 = 1$ since $b \neq b^{-1}$ so $t^2 = b$ and $1 = t^2b$ is not possible. Therefore $\omega_1^{g\phi^i} = \omega_1$ only if $\phi^i = g = 1$. Also, $\omega_1^{gj\phi^i}$ is a pair of two non-squares, so can never equal ω_1, which is a pair of squares. Thus $\{\omega, \omega_1\}$ is a base. □

Proposition B.4. *Let* (G, Γ) *be a group with socle* $\mathrm{PSL}(2, q)$ (*where* $q > 3$) *acting on projective points. Let* $\Omega = \Gamma^{\{2\}}$ *and let* $\omega = \{\alpha, \beta\}$ *lie in* Ω*. Let* $R \subseteq \Omega$ *be a* G_ω*-orbit of largest size. Then* $\#G_{\omega,\lambda} \leqslant 2$ *for* $\lambda \in R$*, and* G_ω *acts faithfully on* R*. If* R *is not contained in the adjacent-point set* $\Lambda_\Gamma(\omega)$ *then the number of points* λ' *of* R *such that* $G_{\omega,\lambda,\lambda'} \neq 1$ *is at most* $2\log_p q$*, where* p *is the prime dividing* q.

Proof. As before, we can identify Γ with the projective line, and take $\alpha = \infty$ and $\beta = 0$. The elements of G can be expressed as a product σ^*h of a transformation h of the form $x \mapsto (ax + b)/(cx + d)$ and a map σ^* induced by the action of an automorphism σ of $\mathbb{F}_q$. The elements of $G_{\alpha,\beta}$ are those for which h is a scalar transformation $x \mapsto lx$ for some $l \in \mathbb{F}_q$.

If $G = \mathrm{PGL}(2, q)$ then the orbits of G_ω on Ω are as follows. There is $\{\omega\}$ and one other orbit, of size $2(q - 1)$, which together make up the adjacent-point set $\Lambda_\Gamma(\omega)$. All the others are $G_{\alpha,\beta}$-orbits of size $q - 1$, of the form $X_k = \{\{a, ak\} \mid a \in \mathbb{F}_q^\times\}$ where the different orbits are obtained by varying the value of k in $\mathbb{F}_q$, except if q is odd then the orbit containing $\{1, -1\}$ has size $(q - 1)/2$. If $G = \mathrm{PSL}(2, q)$ then the orbits are the same, except that the orbit of size $2(q - 1)$ splits into two orbits of size $q - 1$. In both cases G_ω acts regularly on any largest suborbit. In the general case, a largest suborbit is either contained in $\Lambda_\Gamma(\omega)$ (and so has size $q - 1$ or $2(q - 1)$) or is a union of sets X_k each of size $q - 1$.

First we show that if a largest suborbit R is contained in $\Lambda_\Gamma(\omega)$ and has size at least $\frac{1}{2}\#G_\omega$ then G_ω acts faithfully on it. We may assume that G_ω does not act regularly on R, so by the preceding proposition, G acts 3-transitively on Γ and so R must have size $2(q - 1)$. By the assumption on the size of R, the stabilizer of a point of R in G_ω must have size 2, which means that G must be the extension of $\mathrm{PGL}(2, q)$ by a field automorphism of order 2. This field automorphism will fix the pair $\{0, 1\}$, for example, but as there must be elements of $\mathbb{F}_q$ which it does not fix, it cannot fix every pair of the form $\{0, a\}$ where $a \in \mathbb{F}_q$. Thus G_ω acts faithfully on R in this case.

Now we consider the other possibility for R, and show first that there always exists a suborbit of size at least $\frac{1}{2}\#G_\omega$, so any largest suborbit R has size at least $\frac{1}{2}\#G_\omega$. Secondly we will show that if R is a union of sets X_k then G_ω acts faithfully on it, and if G_ω does not act regularly on R then each point stabilizer in the action of G_ω on R fixes at most $2\log_p q$ points of R. The result will then follow.

Let z generate $\mathbb{F}_q$, so z is moved by every non-trivial automorphism of $\mathbb{F}_q$. Consider the action of a general element σ^*h of $G_{\alpha,\beta}$ on $\{1, z\}$; here h is a scalar transformation $x \mapsto lx$

for some $l \in \mathbb{F}_q$. If $\{1, z\}^{\sigma^* h} = \{1, z\}$ then $\{l, lz^\sigma\} = \{1, z\}$. So if either σ^* or h is not the identity we must have $l = z$ and $lz^\sigma = 1$, which gives $z^\sigma = z^{-1}$; this is only possible if $q \leqslant 4$. Thus if $q \neq 4$ then $G_{\alpha,\beta}$ acts semi-regularly on the G_ω-orbit containing $\{1, z\}$, which thus has size at least $\frac{1}{2}\#G_\omega$. However if $q = 4$ then the group of field automorphisms of $\mathbb{F}_q$ has size 2 so if G is 3-transitive then the suborbit contained in $\Lambda_\Gamma(\omega)$ has size at least $\frac{1}{2}\#G_\omega$. If G is not 3-transitive then G_ω acts regularly on a largest suborbit, by the previous proposition.

Now we consider certain special elements, g, g' of G_ω, and show that for each k at least one of these elements fixes a point in X_k, but that neither g nor g' fixes more than two points of any set X_k of size $q - 1$.

In the case $p = 2$, consider the map $g : x \mapsto x^{-1}$, which lies in G_ω. Let $l \in \mathbb{F}_q$ have $l^2 = k^{-1}$ (this is always possible since 2 does not divide the order of $\mathbb{F}_q^{\times}$; however there is only one element l with this property). Then $\{l, lk\}^g = \{l, lk\}$, but this is the unique pair in X_k that is fixed by g, since the only element of $\mathbb{F}_q$ fixed by g is 1 (which equals -1) so any pair $\{a_1, a_2\}$ fixed by g must have $a_1^g = a_2$ and $a_2^g = a_1$.

If p is odd, more care is needed. By the previous proposition, if G_ω does not act regularly on a largest suborbit, then G must act 3-transitively on Γ. Furthermore if G is a 3-transitive extension of $\mathrm{PSL}(2, q)$ by a group of order 2, then G_ω still acts regularly on the suborbit of size $2(q - 1)$ that is contained in $\Lambda_\Gamma(\omega)$. This means that we may actually assume that G contains $\mathrm{PGL}(2, q)$, as the only 3-transitive extension of $\mathrm{PSL}(2, q)$ that does not contain $\mathrm{PGL}(2, q)$ contains $\mathrm{PSL}(2, q)$ as a subgroup of index 2. It follows that g (which has determinant -1 so does not lie in $\mathrm{PSL}(2, q)$) does lie in G. However, there are now either 0 or 2 elements l such that $l^2 = k^{-1}$, depending upon whether k is a square in $\mathbb{F}_q$ or not. Since the only elements of $\mathbb{F}_q$ fixed by g are ± 1, and the pair $\{1, -1\}$ lies in the set X_{-1} that has size $(q - 1)/2$ not $q - 1$, we see that if the set x_k has size $q - 1$ then g must interchange the points of any pair in X_k that it fixes, and so g fixes at most 2 pairs in any X_k of size $q - 1$. For half the sets X_k, that is, those for values of k for which k^{-1} has a square root in $\mathbb{F}_q$, the element g has a fixed point in X_k; we now look for an element g' with similar properties, that fixes a point in all those X_k that do not contain fixed points of g. Assume k has no square root in $\mathbb{F}_q$. Let c be an element of $\mathbb{F}_q$ that is not a square. The transformation $g' : x \mapsto cx^{-1}$ lies in G_ω and there are precisely two elements l for which $l^2 = ck^{-1}$. Then $\{l, lk\}^g = \{l, lk\}$ and g' fixes a point of X_k; as before, g' fixes at most 2 points of any set X_k, since g' does not fix any points of $\mathbb{F}_q$.

Thus for any set X_k, there exists a non-trivial element g or g' of G_ω that fixes a point of X_k but does not fix all of X_k, and so if R is a union of sets X_k and the stabilizer of a point λ in R in G_ω is of size 2, then G_ω must act faithfully on R (using transitivity of G_ω on R in the case where λ itself is not fixed by g or g'). Furthermore, g or g' (as applicable) fixes at most 2 points of each set X_k of size $q - 1$, and so if R is a union of sets X_k and $\lambda \in R$ then there are at most $2 \log_p q$ points $\lambda' \in R$ for which $G_{\omega,\lambda,\lambda'} \neq 1$, since R is a union of at most $\log_p q$ sets X_k, all of which have size $q - 1$.

This is sufficient to prove the proposition. □

The following theorem summarises the foregoing propositions.

Theorem B.5. *Let (G, Γ) lie in $\mathscr{L}'$. Let $\omega \in \Omega = \Gamma^{\{2\}}$ and let R be a largest G_ω-orbit. Let $\lambda \in R$. Then $G_{\omega,\lambda}$ has size at most 2, G_ω acts faithfully on R, and either $R \cup \{\omega\}$ is the adjacent-point set $\Lambda_\Gamma(\omega)$ or there are at most $2 \log_p(m - 1)$ elements λ' of R such that $\{\omega, \lambda, \lambda'\}$ is not a base for G (here p is the prime dividing $m - 1$).*

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press

Appendix C. *GAP script to implement UOP algorithm*

This appendix is available to subscribers to the journal at:
http://www.lms.ac.uk/jcm/1/lms97008/appendixc/.

References

1. L. Babai, G. Cooperman, L. Finkelstein and Á. Seress, 'Nearly linear time algorithms for permutation groups with a small base', *Proceedings 1991 ACM International Symposium on Symbolic and Algebraic Computation* (ed. S. M. Watt, American Mathematical Society, Providence RI, 1991), pp. 200–209. 125, 126, 136

2. William Burnside, *The Theory of Groups of Finite Order* (Cambridge University Press, 1911) 2nd edn. 110

3. Peter J. Cameron and John J. Cannon, 'Fast recognition of doubly transitive groups', *Journal of Symbolic Computation* 12 (1991) 459–474. 119, 132

4. John D. Dixon and Brian Mortimer, *Permutation Groups,* Graduate Texts in Mathematics 163 (Springer-Verlag, New York, 1996). 124, 131, 131, 132, 137, 137, 137, 137, 138, 138, 138, 142, 143

5. W. Feit and J. G. Thompson, 'Solvability of groups of odd order', *Pacific J. Math.* 13 (1963) 775–1029. 110

6. Peter M. Neumann, 'Some algorithms for computing with finite permutation groups', *Proceedings of Groups–St. Andrews 1985* (eds E. F. Robertson and C. M. Campbell), L.M.S. Lecture Note Series 121 (Cambridge University Press, 1987), pp. 59–92. 111

7. Martin Schönert and Ákos Seress, 'Finding blocks of imprimitivity in small-base groups in nearly linear time', *Proceedings 1994 ACM-SIGSAM International Symposium on Symbolic and Algebraic Computation* (1994), pp. 154–157. 124, 124, 124

8. Martin Schönert *et al.*, GAP—*Groups, Algorithms and Programming*. (Lehrstuhl D für Mathematik, Rheinisch Westfälische Technische Hochschule, Aachen, Germany, 1994). 119, 130, 130, 135, 136, 138

9. H. Wielandt, *Finite Permutation Groups* (Academic Press, New York, 1964). 110

Graham R. Sharp sharp@maths.ox.ac.uk

The Queen's College
Oxford OX1 4AW

https://doi.org/10.1112/S1461157000000164 Published online by Cambridge University Press