TRANSACTIONS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 00, Number 0, Pages 000–000
S 0002-9947(XX)0000-0

CONSTRUCTIVE RECOGNITION OF PSL$(2, q)$

M.D.E. CONDER, C.R. LEEDHAM-GREEN, AND E.A. O'BRIEN

ABSTRACT. Existing black box and other algorithms for explicitly recognising groups of Lie type over GF(q) have asymptotic running times which are polynomial in q, whereas the input size involves only $\log q$. This has represented a serious obstruction to the efficient recognition of such groups.

Recently, Brooksbank and Kantor devised new explicit recognition algorithms for classical groups; these run in time that is polynomial in the size of the input, given an oracle that recognises PSL$(2, q)$ explicitly.

The present paper, in conjunction with an earlier paper by the first two authors, provides such an oracle. The earlier paper produced an algorithm for explicitly recognising SL$(2, q)$ in its natural representation in polynomial time, given a discrete logarithm oracle for GF(q). The algorithm presented here takes as input a generating set for a subgroup G of GL(d, F) that is isomorphic modulo scalars to PSL$(2, q)$, where F is a finite field of the same characteristic as GF(q); it returns the natural representation of G modulo scalars. Since a faithful projective representation of PSL$(2, q)$ in cross characteristic, or a faithful permutation representation of this group, is necessarily of size that is polynomial in q rather than in $\log q$, elementary algorithms will recognise PSL$(2, q)$ explicitly in polynomial time in these cases. Given a discrete logarithm oracle for GF(q), our algorithm thus provides the required polynomial time oracle for recognising PSL$(2, q)$ explicitly in the remaining case, namely for representations in the natural characteristic.

This leads to a partial solution of a question posed by Babai and Shalev: if G is a matrix group in characteristic p, determine in polynomial time whether or not $O_p(G)$ is trivial.

1. INTRODUCTION

An *explicit recognition algorithm* for a group H takes as input a group G defined on a generating set X, and decides whether or not G is isomorphic to H; if so, it then also computes the image of X under such an isomorphism and enables computation of the image in G under the inverse isomorphism of a given element of H.

We are concerned with the case where G is given as a matrix group over a finite field.

The black box recognition algorithms for classical groups proposed by Kantor & Seress [22] and the recognition algorithm for SL(d, q) in its natural representation proposed by Celler & Leedham-Green [10] have complexity that is polynomial in the size of the field (and the rank of the group). Classical groups have faithful

Received by the editors March 2003 and, in revised form, October 2003.
1991 *Mathematics Subject Classification.* Primary 20C20, Secondary 20C40.
Key words and phrases. constructive recognition; PSL$(2, q)$.
This work was supported in part by the Marsden Fund of New Zealand via grant UOA124.

representations over the defining field (such as the natural representation) whose degrees are independent of the field size, and so the field size may be exponential in the size of the input. Hence more efficient algorithms are needed to deal with these cases. Recently, new explicit recognition algorithms for groups isomorphic to classical groups have been developed by Brooksbank [5] and Brooksbank & Kantor [6]; these run in time that is polynomial in the size of the input, *provided* the same can be achieved for the groups $\mathrm{PSL}(2,q)$.

The development of an explicit recognition algorithm for $\mathrm{PSL}(2,q)$ is a non-trivial task, even when G is not just isomorphic to, but *is* $\mathrm{PSL}(2,q)$. A central difficulty lies in finding an element of order p, that is to say a transvection, as a word in the given generating set. Conder & Leedham-Green [9] present an explicit recognition algorithm for $\mathrm{SL}(2,q)$ in its natural representation, which runs in time that is polynomial in $\log q$. It relies on a discrete logarithm oracle for $\mathrm{GF}(q)$ — that is, an oracle that will provide, for a given non-zero element μ of $\mathrm{GF}(q)$ and a fixed primitive element a of $\mathrm{GF}(q)$, the unique integer k in the range $1 \leq k < q$ for which $\mu = a^k$. For a description of discrete log algorithms, see [25, Chapter 4].

This reduces the explicit recognition of $\mathrm{PSL}(2,q)$, given by a linear or projective representation in the natural characteristic, to the construction of the natural representation from the given one. In this paper we present such an algorithm which runs in time polynomial in the size of the input.

Let $\tau(d)$ denote the number of factors of d. Our primary result is the following.

Theorem 1.1. *Let G be a subgroup of $\mathrm{GL}(d,F)$ for $d \geq 2$, where F is a finite field of the same characteristic as $\mathrm{GF}(q)$; assume that G is isomorphic modulo scalars to $\mathrm{PSL}(2,q)$. Then there exists a Las Vegas algorithm that constructs an epimorphism from G to $\mathrm{PSL}(2,q)$ at a cost of at most $O(d^5\tau(d))$ field operations.*

In theory we are concerned with a group isomorphic to $\mathrm{PSL}(2,q)$, and which will be defined modulo scalars; but in practice we deal with linear groups, so we have a subgroup G of $\mathrm{GL}(F)$ that is isomorphic, modulo scalars, to $\mathrm{PSL}(2,q)$. It is convenient to replace G by its derived group, so that G is isomorphic to $\mathrm{PSL}(2,q)$ or to $\mathrm{SL}(2,q)$. To replace G with its derived group we use, for example, the polynomial-time Monte Carlo algorithm of [4]. When the algorithm of Theorem 1.1 has run, it provides a simple method of computing the image of an element of G in $\mathrm{PSL}(2,q)$, and this can be applied to any element of the group, not just to an element of the derived group.

Let q be a power of a prime p, and let V be a finite-dimensional vector space over a finite field of characteristic p. In summary, we present an algorithm that takes as input a subset X of the matrix group $\mathrm{GL}(V)$ that generates a group G isomorphic to $\mathrm{SL}(2,q)$ or to $\mathrm{PSL}(2,q)$, and constructs the natural projective representation of G by constructing the image of X under some homomorphism of G onto $\mathrm{PSL}(2,q)$. Our algorithm has three basic steps: the first replaces the given representation with an absolutely irreducible representation over a subfield of $\mathrm{GF}(q)$; the second replaces this absolutely irreducible representation with a tensor irreducible representation over $\mathrm{GF}(q)$; the final step constructs the natural representation from this representation, so that G can be explicitly recognised using the algorithm of [9].

The significance of our algorithm lies in its ability to deal with values of q that are large in terms of the size of the input. Our basic algorithm does not apply for certain small values of q: for example, to the Steinberg representation in odd

characteristic, which is in dimension q. However we can address these cases using more elementary techniques.

A faithful linear or projective representation of $\mathrm{PSL}(2,q)$ in cross characteristic has degree which is polynomial in q [18]; thus we may assume that the input group is a matrix group in its natural characteristic.

In addition to its own intrinsic importance, the algorithm has wider importance for the *matrix recognition project* [19]. This project seeks to reduce the problem of constructive recognition of a matrix group G to the recognition of the composition factors of G. In particular, the constructive recognition of classical groups is of central significance to this project. Brooksbank & Kantor [6] reduce the time required to recognise the classical groups to a function whose dependence on the field size is polynomial in $\log q$, given an oracle which recognises $\mathrm{PSL}(2,q)$ explicitly. Hence this work is a central component of a strategy for the constructive recognition of matrix groups.

The organisation of the paper is as follows. In Section 2 we recall Brauer & Nesbitt's [7] structural characterisation of an absolutely irreducible representation of $\mathrm{SL}(2,q)$ which cannot be written over a smaller field; essentially, this is a tensor product of certain symmetric powers of the natural module. In Section 3 we discuss how to tensor-decompose such a representation of $\mathrm{SL}(2,q)$ to construct a tensor indecomposable factor. In Section 4 we show how to construct the natural module of $\mathrm{SL}(2,q)$ from such a symmetric power. In Section 5 we discuss exceptional cases not covered by the main algorithm. In Section 6 we summarise the explicit recognition algorithm of [9] for $\mathrm{SL}(2,q)$ in its natural representation. In Section 7 we present our complete algorithm and consider its complexity. In Section 8 we discuss how to recognise a linear group isomorphic to $\mathrm{SL}(2,q)$ (as opposed to a projective linear group isomorphic to $\mathrm{PSL}(2,q)$), and how to prove that the input group is in fact isomorphic, as projective linear group, to $\mathrm{PSL}(2,q)$. We also discuss a problem of Babai & Shalev [2]. Finally, we report on the performance of an implementation in MAGMA [8].

2. Irreducible representations of $\mathrm{SL}(2,q)$

Let G be a subgroup of $\mathrm{GL}(d,p^a)$ isomorphic to $\mathrm{SL}(2,q)$ or to $\mathrm{PSL}(2,q)$, where $q=p^e$, and let V denote the corresponding G-module of dimension d over $\mathrm{GF}(p^a)$.

The first aim is to construct from this data a non-trivial absolutely irreducible representation of G that cannot be written over a smaller field.

If V is reducible, then we use the MEATAXE [15], [17] to replace V by a section of V on which G (or possibly G modulo a central subgroup) acts faithfully.

Similarly, if V is irreducible, but not absolutely irreducible, then we rewrite V as an absolutely irreducible G-module in smaller dimension over a larger field [15].

If G, as an absolutely irreducible subgroup of $\mathrm{GL}(d,p^a)$, is conjugate to a subgroup of $\mathrm{GL}(d,p^b)$ for some proper divisor b of a, then we find such a conjugating matrix using the algorithm of [13], and thus write G over $\mathrm{GF}(p^b)$.

We may now apply the following theorem of Brauer & Nesbitt [7, §30].

Theorem 2.1. *Let F be an algebraically closed field of characteristic p, and let V be an irreducible $F[G]$-module for $G=\mathrm{SL}(2,q)$, where $q=p^e$. Then $V \simeq T_1 \otimes T_2 \otimes \cdots \otimes T_t \otimes_{\mathrm{GF}(q)} F$, where T_i is the s_i-fold symmetric power S_i of the natural $\mathrm{GF}(q)[G]$-module M twisted by the f_i th power of the Frobenius map, with $0 \le f_1 < f_2 < \cdots < f_t < e$, and $1 \le s_i < p$ for all i.*

Thus V is defined over $\mathrm{GF}(q)$, and may be defined over a proper subfield K.

Given t and (s_i) and (f_i) it is easy to determine the smallest field K over which V is defined as follows. The group A of automorphisms of $\mathrm{GF}(q)$ is cyclic of order e, generated by the Frobenius automorphism. This group acts on the set of all ordered pairs (s_i, f_i), where $0 \leq f_i < e$ and $1 \leq s_i < p$, the Frobenius automorphism replacing (s_i, f_i) by (s_i, g_i) where $g_i = f_i + 1 \bmod e$. Let B be the subgroup that normalises the set of pairs arising from the above tensor decomposition. Then K is the subfield of $\mathrm{GF}(q)$ corresponding to B in the Galois correspondence; that is to say, K is the subfield centralised by B. In particular, if all s_i are equal, as in the Steinberg representation when $s_i = p-1$ for all i, then K is $\mathrm{GF}(p)$. In practice we have obtained K as the smallest field over which V may be written, but have yet to compute (s_i) and (f_i). If K is a proper subfield of $\mathrm{GF}(q)$ we replace the $K[G]$-module V by the $\mathrm{GF}(q)[G]$-module $V \otimes_K \mathrm{GF}(q)$. Thus the above theorem may be reworded as follows.

Theorem 2.2. *Let K be a finite field of characteristic p, and let V be an absolutely irreducible $K[G]$-module for $G = \mathrm{SL}(2,q)$, where $q = p^e$. Suppose that V cannot be written over a smaller field. Then K is a subfield of $\mathrm{GF}(q)$, and $V \otimes_K \mathrm{GF}(q) \simeq T_1 \otimes T_2 \otimes \cdots \otimes T_t$, where T_i is the s_i-fold symmetric power S_i of the natural $\mathrm{GF}(q)[G]$-module M twisted by the f_i th power of the Frobenius map, with $0 \leq f_1 < f_2 < \cdots < f_t < e$, and $1 \leq s_i < p$ for all i.*

The next step is to decompose this tensor product, by finding one tensor indecomposable factor.

3. A tensor indecomposable representation

The aim of this section is to prove the following result.

Theorem 3.1. *Given a subset of $\mathrm{GL}(d,q)$ that generates a non-trivial tensor decomposable representation of $\mathrm{SL}(2,q)$, a non-trivial tensor indecomposable tensor factor of the given d-dimensional module can be constructed in $O(d^5\tau(d))$ field operations, where $\tau(d)$ is the number of divisors of d.*

We prove this theorem by presenting an algorithm to construct such a tensor indecomposable tensor factor. For odd p, the algorithm does not apply if V has dimension p^e, or if $e = 2$ and V has dimension $p(p-1)$ or $(p-1)^2$. For $p = 2$ and $e = 2$ it fails in dimension 4. We return to these cases in Section 5.

The construction of a tensor indecomposable representation of G relies on the theoretical framework and algorithm developed by Leedham-Green & O'Brien [20, 21] for finding a tensor decomposition of a finite-dimensional module over a finite field, or proving that no non-trivial tensor decomposition of this module exists. Here we summarise briefly the necessary background material, first recalling the concept of equivalence of tensor decompositions.

Definition 3.2. A u-tensor decomposition of V is a linear isomorphism α from $U \otimes W$ onto V, where U and W are vector spaces, with U of dimension u. If $\alpha : U \otimes W \to V$ and $\beta : U' \otimes W' \to V$ are tensor decompositions of V then α and β are *equivalent* if there are linear isomorphisms $\phi : U \to U'$ and $\psi : W \to W'$ such that $\alpha = (\phi \otimes \psi)\beta$. If V is an $F[G]$-module, where F is the underlying field and G is a group, then a u-tensor decomposition of V as $F[G]$-module requires U and W as above to be $F[G]$-modules, and α to be an $F[G]$-isomorphism; and in the definition of equivalence ϕ and ψ are required to be $F[G]$-isomorphisms.

A u-projective geometry on V, where u divides the dimension of V, is a projective geometry where the k-flats are of dimension ku, the join of two flats is their sum, and their meet is their intersection. Thus in a u-tensor decomposition of V as above, the subspaces of V that are the images of subspaces of $U \otimes W$ of the form $U \otimes W_0$, where W_0 runs through the set of subspaces of W, form a u-projective geometry on V. More generally, a u-tensor decomposition of V as $F[G]$-module gives rise as above to a u-projective geometry on V where W_0 runs through the set of $F[G]$-submodules of W. This projective geometry is G-invariant, in that the set of flats is G-invariant. In [20] it was shown that this construction of a u-projective geometry from a tensor decomposition of V as $F[G]$-module sets up a one-to-one correspondence between the set of G-invariant projective geometries on V and the set of equivalence classes of tensor decompositions of V as FG-module. A point in the projective geometry corresponding to a u-tensor decomposition of V has dimension u as a subspace.

In [20] an algorithm FindPoint is presented: given as input a subspace $\mathcal{F}$ of V, it determines whether or not $\mathcal{F}$ is a flat in a G-invariant u-projective geometry on V, and, in the affirmative case, returns the corresponding tensor decomposition of V. In [21] this geometrical approach and some other ideas were exploited to develop a practical algorithm for deciding tensor decomposability. Here we only use one component of FindPoint, namely IsPoint, which decides whether or not the supplied $\mathcal{F}$ is in fact a point; it has complexity $O(d^3)$ field operations.

Hence the problem of finding a tensor decomposition of an FG-module V as $U \otimes W$, where U and W are modules for a covering group of G, is equivalent to constructing a point in one of the two corresponding projective geometries: a subspace of V of the form $u \otimes W$ or $U \otimes w$ for $u \in U \setminus \{0\}$ or $w \in W \setminus \{0\}$.

Assume that we have an irreducible G-module V defined over the field $F = \mathrm{GF}(q)$, where G is isomorphic to $\mathrm{SL}(2,q)$ or to $\mathrm{PSL}(2,q)$. By Theorem 2.2,

$$V \cong T_1 \otimes T_2 \otimes \cdots \otimes T_t, \qquad (*)$$

where T_i is the s_i-fold symmetric power S_i of the natural module M twisted by the f_i th power of the Frobenius map, with $0 \le f_1 < f_2 < \cdots < f_t < e$, and $1 \le s_i < p$ for all i. Hence S_i, which is a space of dimension $s_i + 1$, can be viewed as the set of homogeneous polynomials in x and y of degree s_i over F, with the natural action of $\mathrm{SL}(2,q)$. Accordingly G acts as $\mathrm{SL}(2,q)$ on S_i and T_i for odd s_i, and as $\mathrm{PSL}(2,q)$ for even s_i. In particular, G is isomorphic to $\mathrm{SL}(2,q)$ if the number of odd s_i is odd (or if $p = 2$), and is otherwise isomorphic to $\mathrm{PSL}(2,q)$.

If $p = 2$ then $s_i = 1$ for all i, and $\mathrm{SL}(2,q) = \mathrm{PSL}(2,q)$. If p is odd, we need to consider two cases:

(i) Suppose $q \equiv 3 \bmod 4$. Then the elements of $\mathrm{SL}(2,q)$ of projective order $(q-1)/2$ define two rational conjugacy classes in $\mathrm{PSL}(2,q)$, one defined by the elements of $\mathrm{SL}(2,q)$ of order $q-1$, and the other by the elements of order $(q-1)/2$. If $g \in \mathrm{SL}(2,q)$ has projective order $(q-1)/2$, then either g has order $q-1$ and $-g$ has order $(q-1)/2$, or *vice versa*.

(ii) Suppose $q \equiv 1 \bmod 4$. Then $g \in \mathrm{SL}(2,q)$ has projective order $(q-1)/2$ if and only if g has order $q-1$, so the elements of projective order $(q-1)/2$ define a single rational conjugacy class in $\mathrm{PSL}(2,q)$.

It is easy to find $g \in G$ of order $q-1$ for $p = 2$, or projective order $(q-1)/2$ for odd p, by random search; see Section 7 for further discussion. It follows from

the above analysis that if the eigenvalues of g in the natural representation are $\alpha^{\pm 1}$ then α^2 has order $q-1$ if $p=2$, and order $(q-1)/2$ if q is odd.

Let g be such an element, with eigenvalues $\alpha^{\pm 1}$ in the natural representation. If G is isomorphic to $\mathrm{PSL}(2,q)$, then there is an ambiguity in the sign of α, which disappears since only even powers of α arise in this case. (No attempt will be made to compute α, since this is not needed.)

Next let E denote the multiset of exponents of eigenvalues of g on V, expressed as powers of some fixed primitive element β of $\mathrm{GF}(q)$. The tensor decomposition $(*)$ expresses E in terms of t arithmetic progressions, namely the sets of exponents of eigenvalues of g on the symmetric powers T_i. For all i, the length s_i+1 of the ith arithmetic progression will divide the dimension of V, and will be at most p, and E is the multiset obtained by forming all possible sums constructed by taking one summand from each arithmetic progression.

This gives rise to arithmetic progressions within E itself. If $\alpha=\beta^m$, then m is coprime to $(q-1)/2$, and

$$E \equiv \{m \sum_{i=1}^{t} t_i p^{f_i}\} \mod q-1,$$

where t_i takes each of the s_i+1 values in the set $\{-s_i, -s_i+2, -s_i+4, \ldots, s_i\}$. Now, for some j, let us fix t_i for all $i \neq j$, and let t_j vary. This gives rise to an arithmetic progression of length s_j+1.

Assume that the terms of this arithmetic progression correspond to eigenvalues of multiplicity 1. Then the corresponding eigenvectors span a subspace of V of the form $u \otimes T_j$, and such a space is a point in the corresponding tensor decomposition. The tensor decomposition is now readily constructed, as shown in [20, §3]. In particular, the corresponding action of G on T_j is easily obtained. This may be a projective action, so that an action of $\mathrm{PSL}(2,q)$ may be lifted to an action of $\mathrm{SL}(2,q)$.

How do we construct appropriate arithmetic progressions in E? If E contains multiple entries then we remove all multiple entries from E, and we consider ordered pairs of the remaining elements of E as the first two elements in a possibly appropriate arithmetic progression. If no multiple entries existed in the original E, then we know that suitable arithmetic progressions of length s_i+1 exist for all i, and since the product of the s_i+1 is the dimension d of V, we only consider arithmetic progressions of length a multiple of the largest prime dividing d in the hope that this will reduce the number of arithmetic progressions sent to IsPoint. Also, in this case we may clearly restrict attention to arithmetic progressions that contain some randomly chosen but fixed element of E. If E did originally contain multiple entries then, as we shall see later, $s_i \geq (p-1)/2$ for all i, so in this case we only consider arithmetic progressions of length at least $(p+1)/2$.

These arithmetic progressions are tested in order of increasing length, until we find a tensor decomposition of V, or conclude that none can be found in this way. By testing the arithmetic progressions in order of increasing length we can be certain that when a tensor factor is found it will be tensor indecomposable.

Since $|E|$ is at most d there are at most $d^2\tau(d)$ arithmetic progressions to test, at a cost of $O(d^3)$ field operations each, giving a bound of $O(d^5\tau(d))$ field operations.

There are three outcomes to consider:

(a) a tensor decomposition of V is found;

(b) no tensor decomposition is found, and every entry in E has multiplicity 1, in which case V is tensor indecomposable; or
(c) no tensor decomposition is found, and E has elements of multiplicity greater than 1.

We need to determine what conclusions can be drawn from the last of these possibilities. If V is tensor decomposable then clearly $e > 1$, and we shall see that if $e > 1$ and E has elements of multiplicity greater than 1 then V is tensor decomposable. This is the situation that we now investigate.

The multiset E may be written as

$$\{m \sum_{i=1}^{e} (s_i - 2u_i)p^{i-1} \ : \quad 0 \leq u_i \leq s_i\},$$

where $0 \leq s_i < p$ for each i.

Suppose that some element in E has multiplicity greater than 1, this element arising both from $\{u_i \mid 1 \leq i \leq e\}$ and from $\{u_i' \mid 1 \leq i \leq e\}$. Let $w_i = u_i' - u_i$. Then

$$2 \sum_{i=1}^{e} w_i p^{i-1} \equiv 0 \mod q - 1. \qquad (**)$$

Let S denote the left hand side of (**). We may assume that $S \geq 0$. Since $-p < w_i < p$, we know that if $S = 0$ then $w_i = 0$ for all i, and this does not correspond to a multiple entry in E. Hence we may assume that $S > 0$.

If $S = 2(q-1)$ then $w_i = p - 1$ for all i, which implies that $s_i = p - 1$ for all i since $w_i \leq s_i \leq p - 1$. This means that we have the Steinberg representation. For p odd other multiplicities will also arise from $S = q - 1$, and since no suitable arithmetic progressions can be found, all of these must be treated as exceptional cases. If $p = 2$ then S must be even. Hence the Steinberg representation is the only source of multiplicities for $p = 2$, and there is just one eigenvalue, namely 1, of multiplicity 2. It follows that if $e = 2$ we have an exceptional case, but for $e > 2$ and $p = 2$ it is easy to find a suitable arithmetic progression, and this case can be resolved using the general algorithm.

This leaves the case $S = q - 1$; hence p must be odd. We now need to solve the equation

$$\sum_{i=1}^{e} w_i p^{i-1} \equiv \sum_{i=1}^{e} \frac{p-1}{2} p^{i-1} \mod q - 1,$$

with $-s_i \leq w_i \leq s_i$. The most obvious solution to this equation is to take $w_i = (p-1)/2$ for all i; but other possibilities arise, in that we may add 1 to w_i for $i > 1$ provided we subtract p from w_{i-1}. However, the condition $-s_i \leq w_i \leq s_i$ must be preserved for all i, and since $s_i < p$, this clearly implies that $w_i = \pm(p \pm 1)/2$ for all i.

Now suppose that $s_j < p - 2$ for some j. Take $u_j = \min(s_j, (p-3)/2)$. Then for any w_j as above, of the two integers $u_j \pm w_j$, one is negative and the other exceeds s_j. Thus neither is a candidate for u_j', and hence any eigenvalue with this value of u_j has multiplicity one. Thus fixing u_j at this value, and also fixing u_i for all other values of i with one exception, and letting u_i vary in the range $(0, \ldots, s_i)$ for this exceptional value of i, a suitable arithmetic sequence is obtained.

Now suppose that $e > 2$, and that $s_j = p - 2$ for some value of j. Since the Frobenius automorphism permutes the entries of the tuple $(s_1, \ldots, s_e)$ cyclically,

we may assume that $j = 2$. Now take $u_1 = (p-3)/2$ and $u_2 = (p-1)/2$. If (u_i) gives rise to an eigenvalue of multiplicity greater than 1, the corresponding tuple (w_i) must start $((p-1)/2, (p-1)/2, \ldots)$, or $(-(p+1)/2, (p+1)/2, \ldots)$ or $(-(p+1)/2, -(p-1)/2, \ldots)$. Now $s_2 = p-2$, and we may assume that s_1 is $p-1$ or $p-2$. But whatever the value of s_1, it is easy to see that the tuple $(u_i') = (u_i) \pm (w_i)$ has one of the first two entries out of bounds. That is to say, either $u_i' < 0$ or $u_i' > s_i$. Thus if the first two components of (u_i) are fixed as above, and all but one of the other u_i is fixed (here we use $e > 2$), a suitable arithmetic progression is obtained.

If $e = 2$ and $s_i \geq p-2$ for each value of i again no suitable arithmetic progression can be found, and these cases are also considered in Section 5.

If $s_i < (p-1)/2$ for some i then $u_i < (p-1)/2$; so no multiple eigenvalues can arise in this case, as stated above.

4. Decomposing symmetric powers

We have now reduced the problem of constructing the natural representation of $G = \mathrm{SL}(2, q)$ to the case where the input module V is the s-th symmetric power of the natural module M for some $s < p$, and we wish to construct the natural representation for G. Naturally we may assume that $p > 2$ and $s > 1$.

The aim of this section is to prove the following result.

Theorem 4.1. *Let V be the s-th symmetric power of the natural module for $\mathrm{SL}(2, q)$, where $s < p$. Let $\mathrm{SL}(2, q) = \langle g, h \rangle$ where g has projective order $(q-1)/2$. Given the elements of $\mathrm{GL}(V)$ that define the action of g and h, an ordered basis for V can be constructed that exhibits V as the symmetric power in $O(s^4)$ field operations. Given $k \in \mathrm{GL}(V)$, written with respect to this basis, one can determine in $O(s^2)$ field operations whether or not k lies in the image G of $\mathrm{SL}(2, q)$ in $\mathrm{GL}(V)$. Given that k is the image in $\mathrm{GL}(V)$ of some element u of $\mathrm{SL}(2, q)$, the image of u in $\mathrm{PSL}(2, q)$ can be determined in a fixed number of field operations, independent of q or s.*

Note that the elements g and h could be readily found by random search; in the context in which the theorem is used such elements will be known.

We prove this theorem by presenting an algorithm to accomplish the necessary task. The argument in this section does not cover the following exceptional cases: $q \leq 5$; and $s \geq (p-5)/2$ where $q = p \geq 7$ (unless $p = 13$ and $s = 4$). These exceptional cases – where q is small compared to the size of the input – are considered in Section 5.

It is convenient to abuse notation, and to write g and h for the images of these elements in $\mathrm{GL}(V)$. Let g have eigenvalues $\alpha^{\pm 1}$ in the natural module M, and let x and y be corresponding eigenvectors. Then V can be taken to be the set of homogeneous polynomials in x and y of degree s, and $\{x^s, x^{s-1}y, \ldots, xy^{s-1}, y^s\}$ is then a basis of V consisting of eigenvectors for g. Finding such an ordered basis will exhibit the structure of V, as required. The corresponding eigenvalues can be arranged to form a geometric progression, with common ratio α^2, or, reversing the order, with common ratio α^{-2}.

Our proof relies on two assumptions: the eigenvalues of g on V all have multiplicity 1, and can be arranged in only one way (up to reversal of order) as a geometric series, so that $\alpha^{\pm 2}$ can be determined. This implies $q > 5$. The assumption of uniqueness means that only one possible value for α^2 need be found. We will later show that these assumptions hold for all but the exceptional cases mentioned above.

The choice between α and α^{-1} is a choice between x and y; so suppose that g, acting on the natural module, has eigenvalue $\pm\alpha$ on x and $\pm\alpha^{-1}$ on y. Then the eigenvectors $x^s, x^{s-1}y, \ldots, xy^{s-1}, y^s$ of g acting on V can be constructed, in order, as above, up to non-zero scalar multiples. Choosing these eigenvectors at random, we have an ordered basis

$$(\theta_0 x^s, \theta_1 x^{s-1}y, \ldots, \theta_{s-1}xy^{s-1}, \theta_s y^s)$$

of V where the θ_i are unknown. Since multiplying all the basis elements by the same scalar does not change the matrices, we may assume that $\theta_0 = 1$. Since multiplying y by a scalar divides θ_1 by the same scalar, we may also assume that $\theta_1 = 1$. The other θ_i will be calculated.

The element g was needed only to find the above basis for V. In particular, we do not need to calculate $\alpha^{\pm 2}$, but only decide on its uniqueness. We will return to the question of uniqueness later.

Accordingly, suppose that h as above has matrix

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

in its action on the natural module with respect to the ordered basis (x, y). We are of course not given this matrix. Suppose first that $ab \neq 0$.

The image of the first basis vector $\theta_0 x^s$ under h is

$$\theta_0(ax + by)^s = a^s\theta_0 x^s + sa^{s-1}b(\theta_0/\theta_1)\theta_1 x^{s-1}y + \cdots,$$

from which we can read off the coefficients $a^s, sa^{s-1}b(\theta_0/\theta_1)$, etc., and hence find $(a/b)(\theta_1/\theta_0), (a/b)(\theta_2/\theta_1)$, etc. Since $\theta_0 = \theta_1 = 1$, this determines a/b and all of the remaining θ_i.

Clearly $a = 0$ if and only if x^s is an eigenvector for h, and $b = 0$ if and only if y^s is an eigenvector for h. Thus the condition $ab \neq 0$ can be checked. If $ab = 0$ then $cd \neq 0$, since h neither fixes nor interchanges the eigenspaces of g, as G is generated by $\{g, h\}$. So if $ab = 0$ then the θ_i may be determined by interchanging the roles of x and y. In this way we can construct an ordered basis $(x^s, x^{s-1}y, \ldots, y^s)$ for V, corresponding to some ordered basis (x, y) for M, and take the θ_i to be all equal to 1.

Now let h be an arbitrary element of G, and again denote the matrix of h on M by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

Suppose first that $a \neq 0$; equivalently that x^s is not an eigenvector for h. The image of the first basis vector x^s is now

$$(ax + by)^s = a^s x^s + sa^{s-1}bx^{s-1}y + \cdots.$$

From this we can read off a^s and b/a. The image of the second basis vector $x^{s-1}y$ under h is

$$(ax + by)^{s-1}(cx + dy) = a^{s-1}cx^s + \big((s-1)a^{s-2}bc + a^{s-1}d\big)x^{s-1}y + \cdots.$$

From the first term, we can read off $a^{s-1}c$. Now since we already know a^s, this gives c/a. From the second term, we can read off $(s-1)a^{s-2}bc + a^{s-1}d$. Having computed a^s and bc/a^2, we can compute the first of these two monomials and hence also $a^{s-1}d$ and d/a. We have now computed b/a and c/a and d/a. Thus the image of h in $\mathrm{PSL}(2, q)$ has been computed.

Similarly we can deal with the case $a = 0$ but $d \neq 0$.

If $a = d = 0$ then h maps x^s to $b^s y^s$ and $x^{s-1}y$ to $b^{s-1}cxy^{s-1}$. From this we can read off b/c, which again determines the image of h in $\mathrm{PSL}(2, q)$.

If $a = 0$, this is discovered by observing that h maps the first new basis vector for V to a multiple of the last. Similarly, one can decide whether or not $d = 0$.

Thus, if $h \in G$ is given, then the image of h in $\mathrm{PSL}(2, q)$ can be determined by looking at four entries of the matrix representing the action of h on the given symmetric power of the natural module. Deciding whether or not $h \in G$ from this matrix is of course more expensive. If G is known to act faithfully on this module, so that h is determined by the above matrix, then it is a triviality to extend the above analysis to decide whether or not h lies in G. But in general we will not have this assurance, and will have to proceed as in Section 8.

What of our assumption that the eigenvalues of g have multiplicity one? Multiple eigenvalues occur if and only if $s \geq (q-1)/2$, which implies $q = p$. If for example $q = p$, and $s = p - 3$, then the multiset of eigenvalues consists of the identity with multiplicity 1, and the other quadratic residues with multiplicity 2.

What of our second assumption that the eigenvalues of g on V can be arranged in only one way (up to reversal of order) as a geometric series? Assume that there are no multiple eigenvalues (and so $s \leq (q-3)/2$), but that the set E of exponents of α that arise can be ordered in at least two essentially different ways to form an arithmetic progression. We may take E to be an arithmetic progression with common difference 2, either having an odd number of terms (when s is even) with 0 as the central element, or having an even number of terms with ± 1 as the two central elements (when s is odd). Thus every element of E is even, or every element of E is odd. If E consists of all even residues modulo $q-1$, so that the eigenvalues consist of all the non-zero quadratic residues of $\mathrm{GF}(q)$, then the eigenvalues can be ordered to form a geometric series with common factor β^2 for any primitive $\beta \in \mathrm{GF}(q)$. A similar observation applies if E consists of all the odd residues modulo $q-1$.

If $s = (q-3)/2$, then $q = p$ since $s < p$. It is now easy to see that distinct arithmetic progressions occur for $q = p$ and $s = (q-3)/2$ if and only if $p \geq 7$.

Now suppose that $s < (q-3)/2$. We consider the possibility that multiplication by some integer k prime to $q-1$ should induce a permutation of E modulo $q-1$, where k is not congruent to ± 1 modulo $q-1$ if the elements of E are odd, and $2k$ is not congruent to ± 1 if the elements of E are even. We may take E to be the set of all odd or of all even integers in the interval $[-s, s]$, and so s is the largest element of E.

Suppose first that s is even. Replacing k by $-k$ if necessary, we may assume that $1 < 2k \leq (q-1)/2$, and since k is prime to $q-1$ and we may assume $q > 3$, it follows that $2k < (q-1)/2$. Since $2k \in E$ it follows that $2k \leq s$. Now $s + 2 - 2k \in E$, but $s + 2 \notin E$, so when E is ordered as an arithmetic progression with common difference $2k$, it follows that $s + 2 - 2k$ is the last element of E. If $s + 4 \notin E$ then since $s + 4 - 2k \in E$ we now have a contradiction, as E cannot have two last elements. Thus $s + 4 \in E$ and the only element of the same parity as s missing from E is $s + 2$. Hence we can deduce that E contains exactly $(q-3)/2$ elements, and so $s = (q-5)/2$. Since $(q-3)/2$ is even and $s = (q-5)/2$ is even, $q \equiv 1 \bmod 4$, and E consists of the even residues other than $(q-1)/2$. In this case multiplying by any odd residue modulo $q-1$ maps $(q-1)/2$ to itself, and hence maps E to

itself, and in general we have an exceptional case if $s = (q-5)/2$ for $q \equiv 1 \bmod 4$. If $p = 13$ there is just one arithmetic progression.

Now suppose that s is odd. Then we may assume that $k < (q-1)/2$, and it follows as before that $k < s$ and $s + 4 \in E$. This implies that $q \equiv 3 \bmod 4$, and E consists of the odd residues other than $(q-1)/2$. Again this gives rise in general to an exceptional case.

In summary, the elements of E can be arranged in more than one distinct way as an arithmetic progression if and only if $q = p$ and one of the following holds:

(1) $s = (p-3)/2$ and $p \geq 7$, so that E consists of all even or odd residues; or
(2) $s = (p-5)/2$ and $p \geq 11$, $p \neq 13$, so that E consists of all but one of the even or all but one of the odd residues.

5. Degenerate cases

We have observed that the tensor decomposition algorithm does not apply if p is odd and V has dimension p^e (the Steinberg representation), or if p is odd and $e = 2$ and V has dimension $p(p-1)$ or $(p-1)^2$, or if $p = 2$ and V has dimension 4. Further, if $e = 1$ the symmetric power algorithm does not apply if $s \geq (p-5)/2$ and $p \geq 7$ (unless $p = 13$ and $s = 4$). In these cases q is $O(d)$, so it is easy to recognise G explicitly using more elementary techniques which we now outline.

Omitting the trivial case $p = q = 3$ and $s = 2$, assume $q \geq 4$. Following Theorem 2.2, we reduce to the case where V is an absolutely irreducible $\mathrm{GF}(q)[G]$-module. We explicitly recognise G as follows. Since q is small, we find $x \in G$ of order p by random search. Then x has just one eigenspace on a symmetric power, and this has dimension 1. Hence x has just one eigenspace, say $\langle v \rangle$, on the module in question, as this is a tensor product of symmetric powers. The orbit of this space under the action of G has length $q+1$, and we choose a (permutation group) base for the action consisting of any three distinct spaces in the orbit, say $\{\langle v_i \rangle : 1 \leq i \leq 3\}$, where $v_1 = v$. We find g of order $(q-1)/2$ for q odd, and order $q-1$ for q even, that fixes $\langle v_1 \rangle$ and $\langle v_2 \rangle$, and a of order 2 that interchanges these two spaces. Now x, x^a and g, when referred to an appropriate basis, lift in $\mathrm{SL}(2,q)$ to its Chevalley generators

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \qquad \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \qquad \begin{pmatrix} \alpha & 0 \\ 0 & \alpha^{-1} \end{pmatrix},$$

where α is a primitive element of $\mathrm{GF}(q)^\times$ (or possibly minus a primitive element if $q \equiv 3 \bmod 4$).

We now use this explicit recognition to decompose the tensor product, when $e \geq 2$, and to recognise the symmetric power when $e = 1$. Thus, to find a tensor decomposition of the given G-module V, having recognised G explicitly, we can construct the required tensor product $U \otimes W$, and find an explicit isomorphism between this module and V; we can either do this using the algorithm of Holt & Rees [15] which has complexity $O(d^3)$ field operations, or we can construct one readily by observing that an isomorphism must send the eigenspace of x in $U \otimes W$ to the eigenspace of x in V. There is a certain ambiguity in the choices for U and W when the given module V has dimension $p(p-1)$ and $e = 2$. Then U and W are the $(p-1)$st and $(p-2)$nd symmetric powers of the natural module, up to a choice of a Frobenius automorphism of G. Since $e = 2$, there are precisely two possibilities, and the choice can be settled by trial and error. (This ambiguity is already present in the isomorphism between G and $\mathrm{PSL}(2,q)$.) Since $e = 2$ the

automorphisms of $\mathrm{PSL}(2,q)$ induced by conjugation in $\mathrm{GL}(2,q)$ form a subgroup of index 2 in the full automorphism group, the Frobenius automorphism giving rise to the non-trivial coset. The case of a symmetric power is similar, but simpler. When V has been explicitly recognised as a symmetric power, a basis for V consisting of monomials in x and y of degree s has been found, and so our algorithm for recognising G explicitly can be used.

6. The natural representation

The explicit recognition algorithm for $\mathrm{SL}(2,q)$ in its natural representation takes as input a generating set X for $\mathrm{SL}(2,q)$, and writes any element of $\mathrm{SL}(2,q)$ as a word on X.

A *canonical generating set* for $\mathrm{SL}(2,q)$ in its natural representation consists of a lower-triangular matrix, an upper-triangular matrix and a diagonal matrix

$$\begin{pmatrix} 1 & 0 \\ b & 1 \end{pmatrix}, \qquad \begin{pmatrix} 1 & c \\ 0 & 1 \end{pmatrix}, \quad \text{and} \quad \begin{pmatrix} \alpha & 0 \\ 0 & \alpha^{-1} \end{pmatrix},$$

where α is a primitive $(q-1)$th root of 1 and $b, c \in \mathrm{GF}(q)^\times$. If $b = c = 1$ we refer to this set, and to its image in $\mathrm{PSL}(2,q)$, as a Chevalley generating set. The preprocessing stage of the algorithm finds such a canonical generating set as words on X. A word for an arbitrary element of $\mathrm{SL}(2,q)$ on the canonical generating set is now obtained readily by echelonising the corresponding 2×2 matrix, and so we can write any element as a word on X.

For completeness, we summarise the algorithm of [9] to construct such a canonical generating set for $\mathrm{SL}(2,q)$. Let $H = \langle X \rangle$ be the input group generated by 2×2 matrices of determinant 1.

(1) Find by random search an element A in H of order $q-1$.
(2) Compute eigenvectors u and v of A, with corresponding eigenvalues α and α^{-1}.
(3) Let B be a random conjugate of A.
(4) Find an element C of H and an integer i such that B^iC fixes $\langle u \rangle$. Since A and B^iC lie in $\mathrm{SL}(2,q)$ and have a common eigenvector u, their commutator $S = [A, B^iC]$ is a transvection fixing u.
(5) Similarly, find a random element D of H and a j such that B^jD fixes $\langle v \rangle$ and $T = [A, B^jD]$ is not trivial. Now, T is a non-trivial transvection fixing v.
(6) Write S, T, A with respect to the ordered basis (u, v) to obtain the canonical generating set for $\mathrm{SL}(2,q)$.

Step 4 is the core of the algorithm. Observe that i exists if and only if $\langle u \rangle C^{-1}$ lies in the orbit of $\langle u \rangle$ under $\langle B \rangle$. In odd characteristic, since B has order $q-1$, this orbit consists of approximately half the 1-dimensional subspaces of the natural module; in even characteristic, the orbit has length $q-1$ and consists of every 1-dimensional subspace other than the eigenspaces of B. Hence a suitable C can be found easily. The requirement that B^iC fixes $\langle u \rangle$ is equivalent to the condition that $\alpha^{2i} = \mu$, where $\mu \in \mathrm{GF}(q)$ is determined by the choice of C, which choice must be made in such a way that μ is a square in $\mathrm{GF}(q)$. Computing i relies on a discrete logarithm oracle.

7. The algorithm and its complexity

Assume we are given as input a subset X of $\mathrm{GL}(V)$, where V is a finite dimensional vector space over a field with the same characteristic as $\mathrm{GF}(q)$, and that $G = \langle X \rangle$ is isomorphic to $\mathrm{SL}(2,q)$ or to $\mathrm{PSL}(2,q)$. We consider the complexity of our algorithm to recognise G and discuss membership testing in G.

Following Section 2, we construct an absolutely irreducible representation of G over $\mathrm{GF}(q)$. Following Section 3, we obtain a tensor indecomposable representation of G, and then following Section 4, we construct the natural linear (or projective) representation of G. Following Section 6, we find canonical generators of $\mathrm{PSL}(2,q)$ as words in the images of the elements of X.

The complexity of the MeatAxe is generally $O(d^3)$ field operations, but in one case is $O(d^4)$ [15], [17]. The complexity of writing an irreducible representation in smaller dimension over a larger field is $O(d^3)$ field operations in the smaller field [15]. The complexity of writing an absolutely irreducible representation in the same dimension over a smaller field is $O(d^3)$ field operations in the larger field [13]; this also assumes the existence of a discrete logarithm oracle for the smaller field. Here the smaller field is $\mathrm{GF}(q)$, and we already use a discrete logarithm oracle for this field for the explicit recognition of $\mathrm{PSL}(2,q)$ in the natural representation.

The complexity of constructing a tensor decomposition of the given module, or of proving it to be tensor indecomposable, was given as $O(d^5\tau(d))$ field operations in Section 3. This estimate was based on an analysis of the general case, when the set E of exponents of eigenvalues of multiplicity 1 contains a suitable arithmetic progression. In the exceptional cases, when this does not apply, q is small enough for elementary methods, as in Section 5, to be faster than the above bound. The estimate of $O(d^5\tau(d))$ comes from a factor of d^3, arising from the cost of IsPoint, a factor of d^2 arising from the possible choices of the first two elements of the arithmetic progression, and a factor of $\tau(d)$, the number of divisors of d, as the number of possible lengths of the arithmetic progressions.

This bound is probably significantly too high. It takes no account of the restrictions imposed on the arithmetic progressions that need to be considered. For example, if all eigenvalues have multiplicity 1, the length of the progression is required to be a multiple of the largest prime dividing d, and it is also required to contain some fixed element of E. The former condition should become more effective as the prime in question increases, but is of course vacuous if d is a power of 2. Perhaps a more significant consideration is the fact that many pairs of elements of E will not be the first two terms of any arithmetic progression of length dividing d and passing through a given element of E. Similar considerations might produce a better bound when there are multiple eigenvalues, as in this case we need only consider arithmetic progressions of length at least $(p+1)/2$.

The function $\tau(d)$ is rather ill-behaved, having average order $\log(d)$. (All logarithms are natural.) The normal order of $\log(\tau(d))$ is $\log(2)\log\log(d)$, and if $\epsilon > 0$ then

$$\tau(d) < 2^{(1+\epsilon)\log(d)/\log\log(d)}$$

for all sufficiently large d, while also

$$\tau(d) > 2^{(1-\epsilon)\log(d)/\log\log(d)}$$

for infinitely many values of d; see pp. 64, 359 and 262 of [14] for details. On a more mundane level, $\tau(d) \leq 32$ for $d < 1000$; hence, it is not a significant contributor to the practical cost of the algorithm.

The complexity of the symmetric power algorithm is more straight-forward: the small number of matrix multiplications and linear algebra computations give a complexity of $O(d^3)$ field operations.

A central task at various points in the algorithm is to find an element of order a divisor of $(q-1)$ by random search. The proportion of elements of $\mathrm{PSL}(2,q)$ of the required form is $\left(\frac{1}{2}\phi(q-1)\right)/(q-1)$. Since $\phi(q-1)/(q-1) > 1/\log\log q$ (see [23, §II.8]), we expect to make at most $O(\log\log q)$ random selections to find a suitable element. In practice, random elements are constructed using the algorithm of [12]; after an initial preprocessing stage, the cost of obtaining one is just two matrix multiplications.

Some of the procedures used are Las Vegas, for example the MeatAxe; hence the corresponding complexity is Las Vegas.

The recognition algorithm should be considered as having two stages. The first finds a symmetric power of the natural representation within the given representation, and a basis that exhibits this structure; the second finds the image of $g \in \mathrm{GL}(V)$ in the natural representation of G (or proves that the given element does not lie in G). The first stage is carried out *once*, but the second stage may need to be performed many times.

Suppose then that the first stage has been carried out. This can give rise to a change of basis for the given module as follows. First we choose an irreducible section S/T of V and so take the new basis to contain a basis for S that in turn contains a basis for T. We are now only concerned to change basis elements that lie in S but not in T. This section S/T might not be absolutely irreducible. In this case a change of basis is produced that exhibits S/T as a matrix group over the larger field, elements of the larger field being represented by submatrices. It may now be possible to write this section over a smaller field, in which case a further change of basis will exhibit this fact. The relevant section of the given module has now been written over $\mathrm{GF}(q)$, the elements of $\mathrm{GF}(q)$ appearing either as matrix entries or as blocks. A further change of basis exhibits S/T as a tensor product, so the elements of G are now represented by the Krönecker product of two matrices, one of which gives the action on a symmetric product. A final change of basis exhibits the structure of the symmetric product.

Constructive membership testing can now be readily carried out. If $g \in \mathrm{GL}(V)$ is an element of G, then g must preserve the various structures. If so, we perform the necessary change of basis; now, to obtain the image of g in $\mathrm{PSL}(2,q)$, it suffices to read off 4 of the entries of the matrix giving the action of g on the symmetric product. (These are only determined up to multiplication by a scalar in $\mathrm{GF}(q)$, but only the ratio between entries is significant.) This image is then expressed as a word in the canonical generators of $\mathrm{PSL}(2,q)$, and the word is evaluated in the corresponding elements of G; comparing the resultant element of G with g now decides the membership question. Membership testing is dominated by the cost of evaluating the word, which we shall see in Section 8 is $O(d^3)$ field operations, together with the use of the discrete logarithm oracle.

In practice, the recognition algorithm is most likely to be used while recognising a matrix group of large dimension over a field of modest size. If an irreducible

representation arises that is not absolutely irreducible, this may give rise to a representation of small degree over a large field. Hence, practically, we are not so much concerned with the complexity as a function of d as its dependence on q.

Clearly a precise bound to the complexity of the algorithm should make some reference to the size of the given generating set. If the generating set is large, it should be easy to find a small generating set for G to complete the first stage of the algorithm without considering all of the generators. However, if a proof is required that the given generators do in fact generate $\mathrm{PSL}(2,q)$ modulo scalars, rather than some larger group, then every generator must be run through the second stage.

8. Deciding isomorphism

Recall that our algorithm takes as input a subset X of the matrix group $\mathrm{GL}(V)$ which is known to generate a group G isomorphic to $\mathrm{SL}(2,q)$ or to $\mathrm{PSL}(2,q)$, and constructs the natural projective representation of G by constructing the image of X under some homomorphism of G onto $\mathrm{PSL}(2,q)$.

Suppose now that G is isomorphic to $\mathrm{SL}(2,q)$, and an isomorphism $\theta : G/\zeta(G)$ onto $\mathrm{PSL}(2,q)$ has been constructed. How do we lift θ to an isomorphism from G to $\mathrm{SL}(2,q)$? Clearly this problem only arises when q is odd.

To construct such an isomorphism ϕ, suppose first that θ has been lifted to a map, also denoted by θ, from G to $\mathrm{SL}(2,q)$. Then $\phi(g) \in \{\pm\theta(g)\}$ for all $g \in G$.

If the order of g is odd, then the sign is determined by the fact that $\phi(g)$ is of odd order. Also, if we have an expression for g as a product of commutators (G is perfect) then $\phi(G)$ can be computed by applying θ to this expression, as ambiguities in the sign do not affect the value of a commutator. If the order of g is even, we can multiply g by a random product of commutators, say c, to produce an element $h = gc$ of odd order. Then we can compute $\phi(h)$ and $\phi(c)$, and hence $\phi(g)$. The order of an element of $\mathrm{GL}(d,q)$ can be found using the algorithm of [11]; since we need only learn if the order is even or odd, the associated integer factorisation can be avoided and so the cost is $O(d^3)$ field operations.

The proportion of elements of $\mathrm{PSL}(2,q)$ of odd order is greater than $1/2$. We are here concerned only with the case of odd q. Then if u is the 2-adic value of $q-1$ and v is the 2-adic value of $q+1$ it is easy to see that the proportion is $1/2^u + 1/2^v$; since u or v is 1, this proportion is greater than $1/2$.

A second related problem is the following. How do we verify that the input group G is isomorphic modulo scalars to $\mathrm{PSL}(2,q)$? Having constructed a surjection from G to $\mathrm{PSL}(2,q)$, we can find elements of G that map to the Chevalley generators of $\mathrm{PSL}(2,q)$. We must determine the kernel of this map. To do this, we use the short presentation of Todd [24] for $\mathrm{PSL}(2,q)$ on its Chevalley generators. (Here *short* implies that the presentation has length that is polynomial in the rank of the group and the logarithm of the field size.) We first evaluate these relations on the corresponding elements of G. If these all yield scalars, the next step is to evaluate for all g in a generating set of G the image of g in $\mathrm{PSL}(2,q)$, express this element as a word in the Chevalley generators of $\mathrm{PSL}(2,q)$, evaluate this word on the corresponding elements of G to produce some $h \in G$, and finally check that gh^{-1} is a scalar.

In practice, all words are constructed and stored as *straight-line programs*, because the length of a word in a given generating set constructed in n multiplications and inversions can increase exponentially with n, whereas the length of the

corresponding straight-line program is linear in n. One may intuitively think of a straight-line program for $h \in \langle X \rangle$ simply as an efficiently stored group word on X that evaluates to h. More precisely, a straight-line program is a list of pointers or pairs of pointers. If an item contains a single pointer, the pointer will either point to an element x of X, in which case the item represents the word (x); or it points to an earlier item, in which case it represents the inverse of the word represented by that item. If an item contains an ordered pair of pointers, these pointers will point to earlier elements of the list, and the item represents the product in that order of the words defined by these elements. (For example, the list of words $x, y, x^{-1}, x^{-1}y$ on the set $X = \{x, y\}$ may be represented by the straight-line program $[p_x, p_y, 1, (3, 2)]$.) We allow a third type of item consisting of a pointer to a previous item and an integer. If the previous item represents the group element g, and the integer is n, the item represents g^n. This is significant since such items with large values of n arise naturally, and elements of $\mathrm{GL}(d, q)$ can be raised to high exponents with $O(d^3)$ field operations.

More generally, if G is a matrix group in characteristic p defined by a generating set X, one wishes to determine a composition series for G. Following the strategy of the matrix recognition project [19], we first find a composition series for the given module V, and determine $G/O_p(G)$ as the image of G acting on the direct sum of the composition factors of V. Finding a composition series for $G/O_p(G)$, and recognising the composition factors explicitly, enables one to construct random elements of $O_p(G)$. But can we prove for example that $O_p(G)$ is trivial? To do this, it seems necessary to construct a presentation of each composition factor of G on the generating set found for it. This enables one to construct a presentation for $G/O_p(G)$ on the image of X, and hence to obtain, with proof, a generating set for $O_p(G)$ as a normal subgroup of G.

Can we complete this task in polynomial time? To do this, given a generating set X for a simple projective matrix group H, we need to determine the name of H, construct a canonical generating set for H, find a short presentation for H on this generating set, and write an arbitrary element of H as a straight-line program on this generating set, all in polynomial time.

The critical case is when H is a group of Lie type, with natural characteristic p. We may use the (non-constructive) polynomial-time Las Vegas algorithm of Babai *et al.* [1] to name the group. Given an oracle that recognises $\mathrm{PSL}(2, q)$ explicitly, the algorithms of Brooksbank & Kantor [6] construct Chevalley generating sets for the classical groups in time polynomial in the size of the input. Babai *et al.* [3] and Hulpke & Seress [16] give short presentations of every group of Lie type that is not ${}^2G_2(3^{2m+1})$ on its Chevalley generators.

Hence we can provide a partial answer to the following challenge problem of Babai & Shalev [2]:

> (*p-core problem.*) *The big open problem is to decide whether or not* $O_p(G) = 1$ *where* $O_p(G)$ *is the p-core of* G *(the largest normal p-subgroup) and* G *is a matrix group over* $\mathrm{GF}(p)$ *or more generally, a black box group of characteristic p. Is the p-core problem any easier for matrix groups than for black-box groups of characteristic p?*

They define a group H as a *black box group of characteristic p* if H is a black-box group of some encoding length n, and H is isomorphic to a quotient of a subgroup of $\mathrm{GL}(d, p)$ where $d = \lfloor n/\log p \rfloor$, and p is known. The significance of this definition

is that the composition factors of H are either of Lie type in characteristic p, or have order bounded by some polynomial function of n.

As far as current algorithms are concerned, the $O_p(G)$ problem has complexity that is exponential in the encoding length n in the black-box case, since $\mathrm{GL}(2d, p)$ contains copies of $\mathrm{SL}(2, p^d)$, and it is not known how to find a transvection in $\mathrm{SL}(2, p^d)$, as a black-box group, in sub-exponential time. But in the matrix group context, the problem can now be solved in polynomial time given a discrete logarithm oracle for $\mathrm{GF}(q)$, when the input is a group of Lie type for which the recognition problem can be reduced in polynomial time to an $\mathrm{SL}(2, q)$ oracle.

9. Implementation and performance

An implementation of the complete algorithm is publicly available in Magma. The computations reported in Table 1 were carried out using Magma V2.10 on a Pentium IV 2.8 GHz processor. The input to the algorithm is an absolutely irreducible (projective or linear) representation of $\mathrm{SL}(2, p^e)$ given as a subgroup of $\mathrm{GL}(d, p^k)$. In the column entitled "Factors", we list the degrees of the tensor factors for this representation. In the column entitled "Time", we list the CPU time in seconds needed to recognise the group and to construct the isomorphism between the input representation and the natural representation.

Table 1. Performance of implementation for a sample of groups

d	p	e	Factors	Time
12	5	5	3, 4	0.1
27	3	20	3, 3, 3	4.9
36	5	19	3, 3, 4	276.1
45	7	10	3, 3, 5	9.4
48	5	10	3, 4, 4	18.7
48	19	10	3, 4, 4	50.4
60	5	10	3, 4, 5	22.8
60	19	10	3, 4, 5	65.8
72	7	10	3, 4, 6	46.0
108	5	10	3, 3, 3, 4	249.7
128	2	10	7 factors of degree 2	19.8
128	2	19	7 factors of degree 2	222.5

Acknowledgements

We thank the referee for many useful comments and suggestions on the content and exposition.

References

1. László Babai, William M. Kantor, Peter P. Pálfy, and Ákos Seress, "Black-box recognition of finite simple groups of Lie type by statistics of element orders", *J. Group Theory* **5** (2002), 383–401.

2. László Babai and Aner Shalev, "Recognizing simplicity of black-box groups and the frequency of p-singular elements in affine groups", *Groups and computation*, III (Columbus, OH, 1999), 39–62, Ohio State Univ. Math. Res. Inst. Publ., **8**.
3. L. Babai, A.J. Goodman, W.M. Kantor, E.M. Luks, and P.P. Pálfy, "Short presentations for finite groups", *J. Algebra* **194** (1997), 79–112.
4. László Babai, Gene Cooperman, Larry Finkelstein, Eugene Luks, and Ákos Seress, "Fast Monte Carlo algorithms for permutation groups", 23rd Symposium on the Theory of Computing (New Orleans, LA, 1991). *J. Comput. System Sci.* **50** (1995), 296–308.
5. Peter A. Brooksbank, "Constructive recognition of classical groups in their natural representation", *J. Symbolic Comput.* **35** (2003), 195-239.
6. Peter A. Brooksbank and William M. Kantor, "On constructive recognition of a black box PSL(d, q)", *Groups and computation*, III (Columbus, OH, 1999), 95–111, Ohio State Univ. Math. Res. Inst. Publ., **8**, de Gruyter, Berlin, 2001.
7. R. Brauer and C. Nesbitt, "On the modular characters of groups", *Ann. of Math.* **42**, 556–590, 1941.
8. Wieb Bosma, John Cannon, and Catherine Playoust, "The MAGMA algebra system I: The user language", *J. Symbolic Comput.*, **24**, 235–265, 1997.
9. Marston Conder and Charles R. Leedham-Green, "Fast recognition of classical groups over large fields", *Groups and computation*, III (Columbus, OH, 1999), 113–121, Ohio State Univ. Math. Res. Inst. Publ., **8**, de Gruyter, Berlin, 2001.
10. F. Celler and C.R. Leedham-Green, "A constructive recognition algorithm for the special linear group", *The atlas of finite groups: ten years on* (Birmingham, 1995), 11–26, London Math. Soc. Lecture Note Ser., **249**, Cambridge Univ. Press, Cambridge, 1998.
11. Frank Celler and C.R. Leedham-Green, "Calculating the order of an invertible matrix", *Groups and Computation II, Amer. Math. Soc. DIMACS Series* **28**, 55–60. (DIMACS, 1995), 1997.
12. Frank Celler, C.R. Leedham-Green, Scott H. Murray, Alice C. Niemeyer, and E.A. O'Brien, "Generating random elements of a finite group", *Comm. Algebra* **23**, 4931–4948, 1995.
13. S.P. Glasby and R.B. Howlett, "Writing representations over minimal fields", *Comm. Algebra*, **25**, 1703–1711, 1997.
14. G.H. Hardy and E.M. Wright. *An Introduction to the Theory of Numbers.* Fourth edition. OUP. Oxford.
15. Derek F. Holt and Sarah Rees, "Testing modules for irreducibility", *J. Austral. Math. Soc. Ser. A*, **57**, 1–16, 1994.
16. Alexander Hulpke and Ákos Seress, "Short presentations for three-dimensional unitary groups", J. Algebra **245** (2001), 719–729.
17. Gábor Ivanyos and Klaus Lux, "Treating the exceptional cases of the MeatAxe", *Experiment. Math.* **9** (2000), 373–381.
18. Vicente Landazuri and Gary M. Seitz. "On the minimal degrees of projective representations of the finite Chevalley groups", *J. Algebra* **32**, 418–443, 1974.
19. Charles R. Leedham-Green, "The computational matrix group project", in *Groups and Computation*, III (Columbus, OH, 1999), 229–247, Ohio State Univ. Math. Res. Inst. Publ., **8**, de Gruyter, Berlin, 2001.
20. C.R. Leedham-Green and E.A. O'Brien, "Tensor products are projective geometries", *J. Algebra*, **189**, 514–528, 1997.
21. C.R. Leedham-Green and E.A. O'Brien, "Recognising tensor products of matrix groups", *Internat. J. Algebra Comput.*, **7**, 541–559, 1997.
22. William M. Kantor and Ákos Seress. *Black box classical groups.* Mem. Amer. Math. Soc. **149**, 2001.
23. D.S. Mitrinović, J. Sándor and B. Crstici, *Handbook of Number Theory, Mathematics and Its Applications* **351**, Kluwer Academic Publishers, 1996.
24. J.A. Todd, "A second note on the linear fractional group", *Proc. London Math. Soc.* **11**, 1936, 103–107.
25. Igor E. Shparlinski, *Finite fields: theory and computation. The meeting point of number theory, computer science, coding theory and cryptography.* Mathematics and its Applications, **477**. Kluwer Academic Publishers, Dordrecht, 1999.

Department of Mathematics, University of Auckland, Private Bag 92019, Auckland, New Zealand
E-mail address: conder@math.auckland.ac.nz

School of Mathematical Sciences, Queen Mary, University of London, London E1 4NS, United Kingdom
E-mail address: C.R.Leedham-Green@qmul.ac.uk

Department of Mathematics, University of Auckland, Private Bag 92019, Auckland, New Zealand
E-mail address: obrien@math.auckland.ac.nz